



CONTENT OPERATIONS

FROM START TO SCALE

PERSPECTIVES FROM INDUSTRY EXPERTS

EDITED BY CARLOS EVIA

Content Operations from Start to Scale

Perspectives from Industry Experts

Edited by
Carlos Evia



This book was peer-reviewed by two expert practitioners prior to publication
by Virginia Tech Publishing.

Copyright © 2024 Carlos Evia
Individual chapters © 2024 respective authors

First published 2024 by Virginia Tech Publishing for the Virginia Tech
Academy of Transdisciplinary Studies

Virginia Tech Publishing
University Libraries at Virginia Tech 560 Drillfield Drive
Blacksburg, VA 24061

Virginia Tech Academy of Transdisciplinary Studies
College of Liberal Arts and Human Sciences
Liberal Arts and Human Sciences Building
200 Stanger Street
Blacksburg, VA 24061

This work is licensed under the Creative Commons Attribution 4.0
International License (unless stated otherwise within the content of the work).
To view a copy of this license, visit [http://creativecommons.org/licenses
/by/4.0/](http://creativecommons.org/licenses/by/4.0/) or send a letter to Creative Commons, 444 Castro Street, Suite 900,
Mountain View, California, 94041, USA. This license allows for copying
any part of the work for personal and commercial use, providing author
attribution is clearly stated.

Cataloging-in-Publication Data
Content Operations from Start to Scale: Perspectives from Industry Experts /
edited by Carlos Evia.
Includes bibliographical references.

ISBN (PDF): 978-1-957213-87-3
ISBN (EPUB): 978-1-957213-74-3

Suggested citation:

Evia, C. (ed.) 2024. *Content Operations from Start to Scale: Perspectives from Industry Experts*. Blacksburg: Virginia Tech Publishing. DOI: https://doi.org/10.21061/content_operations_evia. License: CC-BY

To read the free, open access version of this book online,
visit https://doi.org/10.21061/content_operations_evia
or scan this QR code with your mobile device:



Contents

List of Figures and Tables	v
Acknowledgments	vii
Foreword	
<i>Deane Barker</i>	ix
Introduction	
<i>Carlos Evia</i>	1
Chapter 1: Defining Content Operations	
<i>Rahel Bailie</i>	13
Chapter 2: The Business Case for Content Operations	
<i>Sarah O'Keefe</i>	25
Chapter 3: Governance for Content Operations	
<i>Kate Kenyon</i>	33
Chapter 4: Operationalizing Content Creation	
<i>Rahel Bailie and Carlos Evia</i>	43
Chapter 5: Customer Experience and Content Operations	
<i>Kevin P. Nichols</i>	75
Chapter 6: Personalization and Content Operations	
<i>Jeffrey MacIntyre</i>	91
Chapter 7: Localization and Content Operations	
<i>Loy Searle</i>	109
Chapter 8: The Technology that Supports Content Operations	
<i>Patrick Bosek</i>	125

Epilogue: You're Here because of People

Jonathan McFadden

147

Afterword

Jason Swarts

155

Contributors

163

Index

167

List of Figures and Tables

Figures

1.1	T-ledger for double-entry bookkeeping	18
1.2	Office clerks in Victorian times	19
4.1	Sample of conventions in the screenwriting format	49
4.2	Sample of an advertisement breaking conventions	50
4.3	Content outputs from a single source of content modules	58
5.1	Representation of a customer content omnichannel experience	76
5.2	Sample customer journey map based on a fictional online grocery store	79
5.3	Task-based customer journey with content mapping	84
6.1	A framework that specifies the mechanics behind Slack notifications	98
6.2	<i>Perso-fail</i> example from a Twitter user complaining about Amazon's multiple offers for toilet seats	103
7.1	Common localization industry definitions	111
7.2	Best practice translation process	120
8.1	Conceptual primary components of a ContentOps supply chain	130
8.2	Base ContentOps infrastructure by supply chain component	131
8.3	ContentOps infrastructure with common connected systems	132

Tables

1.1	A few common content genres, with current approaches and more efficient solutions	21
-----	---	----

3.1	Possible checkpoints for governance review in a ContentOps workflow	37
4.1	Example of an information asset with typical conventions applied	48
4.2	Comparison of elements from two content types	52
4.3	Potential content assets, by stage in the user journey, needed for a single product or service	56
4.4	Potential content variants needed for a single product or service	56
6.1	The four enduring arguments for personalization	92
8.1	Sample battle card for selecting and implementing technology related to ContentOps	126

Acknowledgments

This edited collection was conceived in the spring of 2019 as the fourth edition of *Technical Writing 101*, a book that was originally published in three editions by Scriptorium and written by Alan Pringle and Sarah O'Keefe. Since those initial days of this project, Sarah has remained a key player despite the COVID-19 pandemic and many changes to the book's structure and content. Sarah and many of the current and former staff members at Scriptorium deserve credit for significant work behind this volume. Above all, the contributors who generously wrote chapters for an open-access publication deserve recognition for responding to my pitch of "this book won't make you rich, but it will help explain the concept of Content Operations in industry and academia." Patrick Bosek was a very important force and moved the project forward. Rahel Bailie is a kind and knowledgeable colleague who became an excellent coauthor and gave me more than one content master class during our many Zoom meetings. At Virginia Tech Publishing, Peter Potter, Anita Walz, and Corinne Guimont championed the project, with patience, for years. Rebekka Andersen, Alan Houser, Jimmy Ivory, Tim Lockridge, and Kelly Pender helped me with quick answers, opinions, and overall encouragement during the editing process. This collection would not exist without the support and care of Jane Robertson Evia and Sofia Evia.

Foreword

Deane Barker

I was walking down Wall Street in Lower Manhattan one morning, on my way to a meeting.

The New York Stock Exchange is at the corner of Wall and Broad. As I passed by, in the middle of the street, a film crew was setting up for an “impromptu” interview segment. The gaffers were wrangling lighting equipment, a distinguished middle-aged man—probably a hedge fund manager or something—was in the makeup chair, and another man was putting out cones to keep the area behind the shot clear.

I went into a coffee shop, waited in line too long, and, as I exited, I noticed they were *still* setting up for this interview. I imagine it was a 30-second spot, but they probably spent an hour getting everything ready. It was odd to see the interview from a distance, where it looked fairly unremarkable, and to know that what would make it onto TV screens would be visually perfect.

And that’s when it occurred to me that TV producers lie for a living.

This hedge fund manager didn’t have a perfect complexion. All the buildings made the lighting on this corner pretty awful. And without blocking off the street, there would be dozens of random people walking around in the background.

But a producer has to *produce*. They have to orchestrate and manipulate the world so that it looks the way they need it to. They spend an hour preparing to lie for 30 seconds.

If you’re reading this book, you probably lie for a living too.

You create and manage content. But the world doesn’t know much about the “create and manage” part. All the world knows is that wonderful content

How to cite this book chapter:

Barker, D. 2024. Foreword. In: Evia, C. (ed.) *Content Operations from Start to Scale: Perspectives from Industry Experts*. Pp. ix-xi. Blacksburg: Virginia Tech Publishing.
DOI: https://doi.org/10.21061/content_operations_evia_1. License: CC BY

magically appears. People don't know what it takes to make it so, and it's both your professional responsibility and your fervent hope that they never find out.

Consider that one of the secrets of Google Docs is that it tracks and records every keystroke. So a Google document isn't just a snapshot of text at a moment in time—it's actually a history of every keystroke that it took to get the document to its current state. Theoretically, you could "rewind" a document through every single moment of its creation.

If you're a writer, you're probably horrified by this, and I don't blame you. As a content creator, the imprecise, messy, and sometimes cringe-worthy process it took to get from blank page to finished product is not something you want people to see. You make embarrassing mistakes, you go down blind alleys, you ruthlessly kill your darlings, and you'd just rather the world not know about this, thank you very much.

Content Operations is all of this, writ large.

Like it or not, behind every amazing piece of content is a long, messy chain of humans, processes, events, and tasks that bring it all together. Just like our 30-second interview segment, which took a dozen people working for an hour, there's a lot of behind-the-scenes work that goes into creating a content engine that generates assets your organization can use, keeps them safe and secure, and produces artifacts from them in a reliable and repeatable process.

I've been working in the content technology world for almost 30 years now. The way we create and manage content has certainly evolved and improved, but there's a "floor" in how refined the process can ever become. It's your job to get as close to this floor as you can.

Organizations have an unfortunate tendency to throw technology at this problem, seeking the mythical nirvana of "digital transformation" (what does that even mean?) while ignoring the very real human processes behind the curtain.

You see, one of the beautiful things about content is that it's a time-shifted connection between creator and consumer—two minds coming into a union through content. This means there are humans on both sides of the transaction. While we put a shocking amount of time into ensuring that our content is *consumed* as efficiently as possible, we don't seem to put the same amount of time into ensuring that it's *created* with the same level of efficiency.

This is like a retail goods manufacturer having the most effective sales process in the world but left without anything to sell because no one bothered to make sure the factory was turning out goods. In our case, the consumer is king; the editors ... well, they'll just figure it out, right?

To understand and refine Content Operations, you need to:

- Know how to define and discuss it
- Be able to articulate the business value of it
- Codify your constraints, rules, and standards into a repeatable framework
- Create and enforce process flows that are rigid enough to guide but flexible enough to adapt

- Understand the technology environment you and your editors will need to operate in
- Make tactical decisions about how to adapt your content to your audience's behavior, culture, and needs
- Effectively manage the sometimes fragile humans who make the entire process work

Content Operations isn't for the faint of heart. It's often easier to be a cog in the machine than it is to run the machine. The creative process isn't always clean and precise, and it can probably never be so. It embodies mistakes, destruction, and turmoil. In that sense, it's a snapshot of the human condition.

But content creation and operations is a noble cause. We are the silent servants, standing in the gap, connecting humans across space and time. We delight and entertain. We inform and educate. We motivate and persuade. We light up parts of the brain that might otherwise remain dormant.

And, yes, by abstracting away the messiness and making the finished product appear effortlessly, like magic, we basically tell lies for a living.

But know that they are grand, beautiful lies. And they make the world a better place.

Godspeed and good luck.

Deane Barker
Global Director of Content Management
Optimizely
May 2022

Introduction

Carlos Evia

Content Operations (ContentOps) is the implementation of a strategy that incorporates people, processes, and technologies to optimize content production, so that organizations can leverage their content as business assets while retaining content quality.

—Rahel Bailie

I have been a college professor, working in the disciplinary mixture that we now call Content Operations or ContentOps, since 2004. Before that, as a doctoral student of technical communication at Texas Tech University, I was fortunate to take classes with faculty who, in the early 2000s, inspired me to think about ways to structure, automate, and improve workflows for developing and delivering digital information. Some of those professors were committed to exploring collaborations between the practitioner and academic sides of the discipline. In my early years as a college professor at Virginia Tech, I inherited a course about technical content titled Creating User Documentation. I taught a section of that course in the spring of 2022, and some of the principles in Creating User Documentation have not changed much in 18 years. After all, fundamentals of technical writing are still important, practitioners and academics should collaborate, and audience and task analysis are necessary steps for any deliverable related to technical content. But many other components of the course have evolved as trends and tendencies in industry and academia adapt to incorporate new technologies and processes.

How to cite this book chapter:

Evia, C. 2024. Introduction. In: Evia, C. (ed.) *Content Operations from Start to Scale: Perspectives from Industry Experts*. Pp. 1–12. Blacksburg: Virginia Tech Publishing.
DOI: https://doi.org/10.21061/content_operations_evia_2. License: CC BY

Over the years, one of the main resources for updating the content of college courses (mine and those of some of my professors in graduate school), and for trying to keep up with the fast pace of change in the content disciplines, has been the guest lecture. I have never hesitated to invite colleagues from industry or other universities to visit my class and talk to the students about new ideas or topics that my syllabus was not covering. At the same time, I also did my share of guest lecturing, offering an academic perspective within corporate environments and industry conferences. It is through that mechanism of guest lectures that I met Sarah O’Keefe. Back in 2009, I was looking for an industry-based speaker to visit my class and talk about the importance of structured authoring for technical content. After seeing Sarah give a great presentation on this topic at the conference of the Society for Technical Communication, I contacted her, and she was generous enough to drive from Raleigh, North Carolina, to Blacksburg, Virginia, to talk to a class of 22 students in exchange for lunch and dinner.

Sarah and I stayed in touch and collaborated on a couple of projects, but since 2013 I had been telling her that I wanted to work with her on a new edition of *Technical Writing 101*—a book that she had coauthored with Alan Pringle. I used the three published editions of *Technical Writing 101* as a required textbook in Creating User Documentation for many years, and even today I am including some parts of its template for a documentation plan in my course. The book’s third (and probably last) edition was published in 2009 (Pringle & O’Keefe, 2009), but many of us in the technical communication discipline consider it a classic. The promised fourth edition of *Technical Writing 101*—to be coauthored by Sarah, Alan, and me as an open-access educational resource—became an ongoing conversation when I talked to Sarah at conferences or on Twitter.

In 2019, I had a couple of meetings with Sarah at a conference of the Center for Information Development Management in North Carolina, and we finally started drafting a fourth edition of *Technical Writing 101*. Nevertheless, the writing process hit a major hiccup very early on. In the introduction to the book’s third edition, Sarah and Alan had included the following anecdote:

For technical writers, answering the obligatory “What do you do for a living?” question at a party can have many effects. It can:

- Create more questions: “What’s that?”
- Draw blank stares
- Provoke some minor hostilities: “Did you write that worthless online help that came with my word processing software?” (Pringle & O’Keefe, 2009, p. 19)

As Sarah, Alan, and I developed a plan to revise and update the book, we realized that we just couldn’t write it. Although we still *did* some technical writing (Sarah and Alan as consultants, and I in a small portion of my teaching

and research), we just no longer identified as technical writers. The “What do you do for a living?” question hit us hard. We attempted to answer it with “professional communicator,” “content strategist,” “content developer,” “content specialist,” even “content designer”—until we noticed that the term “content” was unavoidable in any of our most accurate answers, which eventually led to “content operations specialist.” In the very broad terms that we drafted during the conference’s lunch breaks, Sarah and I saw the role of the content operations specialist as a coordinator of the processes (manual and automated) to implement a strategy for developing, deploying, and delivering user-facing content in an organization.

Around that time, Patrick Bosek—who, incidentally, also attended that conference in North Carolina—published on LinkedIn a post attempting to define Content Operations. Patrick’s early definition of the term is as follows: “Content Operations—ContentOps—is the infrastructure that maximizes your content creators’ efforts and guards against procedural errors by automating as much of the content development process as possible” (Bosek 2020).

Other industry colleagues, including Rahel Anne Bailie—who had served with me on a committee in the late 2010s—replied to Patrick’s post and provided their own definitions, which gave me an idea for the book that I still planned to write with Sarah. What if, instead of an introduction to technical writing, we were to write an introduction to ContentOps? What if we brought back the guest lecture model and invited experts from industry to provide chapters that could guide discussion about this topic? Sarah and I invited Patrick and Rahel (who very kindly joined us in many Zoom calls at 9:00 pm—her time—during the height of the COVID-19 pandemic) to join the book project.

That early idea for this book had us writing from the intersection of technical communication and ContentOps. However, we noticed that focusing only on the ContentOps of technical communication would leave out many important conversations and voices from marketing, user experience research, localization, and other disciplines that were also talking about ContentOps. As a result, we expanded the team and invited more industry experts to contribute their “guest lecture” chapters to cover a broad spectrum of professional and enterprise content. It helped that around that time, Rahel started hosting a webinar series for The Content Wrangler network—titled Let’s Talk Content Ops—that expanded our list of industry experts. Kevin Nichols came along, and he brought Loy Searle. Rahel and Sarah invited Jeff MacIntyre. Rahel suggested Kate Kenyon, and I reached out to Jonathan McFadden and Jason Swarts—a kindred-spirit college professor interested in the content professions. Anita Walz and Peter Potter, at Virginia Tech Publishing, patiently waited while I pitched them a different book with a revamped list of contributors every semester. The generous contributors to this edited collection agreed with our original goal of publishing it as an open-access resource.

If you are a practitioner or executive reader and the guest lecture genre does not work for you, think of the contributions compiled in this book as expert presentations inspired by the Talks at Google series or the toolbox talks with specialists at a work site during the lunch break. We aim to provide an introduction to ContentOps, explaining its place in the organization and how we can use its principles to deliver better content at scale. This collection of essays has the purpose of making ContentOps teachable across silos and departments. It offers a convergent approach for industry-based training and college-level courses representing an inclusive definition of professional content. We hope that the chapters in this collection inform decision makers in industry and inspire research and teaching ideas in academia. Above all, we hope that it minimizes the number of blank stares when the term “Content Operations” or “ContentOps” is mentioned in the workplace or the classroom.

But ... what is content?

Spoiler alert: Although this book includes a whole chapter titled “Defining Content Operations,” the definition that Rahel includes in that chapter—and that I included at the beginning of this introduction—is as follows: “Content Operations (ContentOps) is the implementation of a strategy that incorporates people, processes, and technologies to optimize content production, so that organizations can leverage their content as business assets while retaining content quality.”

Throughout this book, our industry experts share Rahel’s vision for the term. We implement the definition of ContentOps as the operationalization of a content strategy. Another spoiler: Rahel defines content strategy as the formation of a plan that codifies a repeatable system that governs the management of content throughout the entire content life cycle. Then again, she has been using a version of that definition since the early 2010s (see, for example, Bailie 2014).

Later in this introduction, I will take a look at how some other authors and organizations are using the term “Content Operations.” Regardless of the specific meaning that an author gives to ContentOps, we need to start by defining what we mean by content.

Some definitions of content are broad. Probably one of the most generous definitions of the term comes from Mike Atherton and Carrie Hane, who see content as “substantive information that is expressed through a medium” (2018, p. 33). Other long-standing definitions focus on the needs of consumers of digital communication products. Kristina Halvorson and Melissa Rach, for example, provide the following definition in their seminal book *Content Strategy for the Web*: “Content is what the user came to read, learn, see, or experience. From a business perspective, the content is the critical information the website, application, intranet, or any other delivery vehicle was created to contain or communicate” (Halvorson & Rach, 2012).

Deane Barker provides a definition that focuses on process and consumption. According to Barker, content is “information produced through editorial

process and ultimately intended for human consumption via publication” (2016a, p. 5). Writing from a disciplinary platform in marketing, the late James Mathewson and his coauthor Mike Moran found their definitions in the eyes and fingers of users, describing content as “what users of digital media consume. It’s what they look for when they search for stuff. It’s what they promote when they like and share items” (2016, p. 13).

After Rahel read an early draft of this introduction, she commented that she uses the following definitions to distinguish terms related to content:

- Copy = *editorial*; texts: meant for human comprehension
- Metadata = *technical*; schemas: meant for machine “understanding”
- Content = *editorial + technical*.

The combination of *editorial* and *technical* to create the kind of content that we are talking about in this edited collection has strong roots in the discipline of technical communication—particularly as it applies to content management practices and tools. In their article “The Current State of Component Content Management: An Integrative Literature Review,” Rebekka Andersen¹ and Tatiana Batova report on a meticulous exploration of academic and industry publications that define the term “content.” Their findings show that content “is used widely to refer to the product- and service-related information that organizations produce for their customers and other stakeholders” (2015, p. 253). Andersen and Batova acknowledge that *content* has not been extensively defined, writing, “the literature tends to focus on definitions of content components, the stand-alone, small units of text used to build information products, such as online Help, user manuals, and training materials” (p. 253). Some of the definitions they quote include the following from Scott Abel: “Any text, image, video, decoration, or user-consumable elements” that help people understand “an organization’s products or services, stories, and brand” (2014, p. 12). Abel expands that definition by pointing out that content can be described in the following ways:

- Contextualized data
- The stuff inside a container
- An extension of the user experience
- A business asset (Abel, 2014, p. 12)

Andersen and Batova also highlight a definition from Joe Gollner, who situates content between data and information in the data-information-knowledge

¹ Rebekka Andersen is my most frequent collaborator in academia. I acknowledge and appreciate her behind-the-scenes input into this introduction (and the whole edited collection).

hierarchy. Gollner's definition views content as "potential information," as the "connective tissue that connects information transactions to data resources and accumulates these transactions into the social construction of grounded knowledge that, in turn, serves as the basis for sound judgment and effective action" (Gollner, 2014). Andersen and Batova started their exploration from their primary disciplinary identity as technical communicators, but they identified a trend that expanded the use of content (and some related terms, like "component content management") "for all product- and service-focused content, including marketing, learning, and technical content, as well as web content" (2015, p. 254).

From a wider perspective, grounded in writing studies, Lisa Dush proposes the following definition: "Content is writing—or composed texts—also conceived of as digital assets, conditional in their shape and value, that are assembled within and pushed out to networks, where human and machine audiences will assess them, assign value to them, consume them, appropriate and repurpose them, extract from them, and push them into other networks. Said differently, as a set of characteristics, content is conditional, computable, networked, and it is—or will be—commodified" (2015, p. 178). That commodification, Dush claims, "*always happens to texts in circulation*" (p. 178; her emphasis). She adds that it "is important to hold this commodification at the core of content's definition, as it encourages healthy suspicion and critical examination" (p. 178).

Quite a few authors are more than healthily suspicious about the commodification of writing-as-content. Miles Kimball has been one of the most vocal academics raising concerns about the commodification of content—specifically in the field of technical writing. Kimball argues that document production that separates form from content has become the domain of two groups: "one smallish group of people paid well to think strategically about design, rhetoric, visualization, presentation, information architecture, technology, and usability, and another much larger, less-well-paid group of people who write fragmented paragraphs that they save to a database, never knowing exactly where or how they will be used" (2016, p. 10). This larger group of writers contributes to the "content machine" that those in more strategic roles have designed. In Kimball's opinion, "this movement toward separating content writing from strategic design makes technical writing of this type look less like a profession and more like a modular, relatively mechanical activity—just like any part of a machine" (p. 9).

The idea of the content machine (which Sarah O'Keefe and Kate Kenyon address in their chapters included in this volume) is a process-focused metaphor for describing ContentOps. Like Alex Masycheff, a content automation specialist in industry who has also addressed the commodification of content (2018), we believe that focusing on user experience and human—and humanistic—endpoints is a more comprehensive and (why not?) optimistic approach

to authoring, processing, and delivering content. Focusing on human activities allows us to maintain a healthy suspicion of abusive content practices, but it also allows us to prepare professionals to avoid those practices and keep humans in the loop.

It would be helpful to describe the commodification of content as the treatment of *some* communication products as business assets. ContentOps—or *effective* ContentOps—does not aim to commodify all communications in an organization. The goals, instead, should be to perceive and treat valuable pieces of content as business assets and to establish and maintain workflows that implement a strategy for the automation of routine tasks related to content creation and processing.

Furthermore, no one is calling for a commodification of all creative content. If one day I finally decide to write the next great Mexican American novel, I probably will not write it in “fragmented paragraphs that [I] save to a database, never knowing exactly where or how they will be used,” unless the novel is about the content machine—which is actually not a bad idea for a novel. It is not as if the automation of business content has the purpose of shattering the *aura*²—in the sense intended by the philosopher and cultural critic Walter Benjamin—of works of art. Rather, it seeks to make consumable content easier and faster to produce and to offer a consistent user experience to its receivers.

The content machine of automation is already here and is giving us personalized experiences, for good or ill. It is up to us to decide whether we let the developers and algorithms control it or if we want to maintain the humanity of ContentOps professionals as one of its key components. Establishing a contrast between working in theory and practice, the chapters included in this collection present experiences and recommendations from content professionals representing technical documentation, marketing, content strategy, and design, who have influenced and conquered the content machine for good.

The purpose of this volume is to present an approach to ContentOps that prioritizes people over tools; cares about diversity, equity, and inclusion in authors and readers; establishes and updates a governance plan built on ethical considerations and effective user experience research; aligns with organizational business goals without ignoring the human beings involved in its processes; implements and revises the appropriate technology tools for responsible automation and personalization of content; and adapts to the global and local needs of diverse audiences.

² “What, then, is the aura? A strange tissue of space and time: the unique apparition of a distance, however near it may be. To follow with the eye—while resting on a summer afternoon—a mountain range on the horizon or a branch that casts its shadow on the beholder is to breathe the aura of those mountains, of that branch” (Benjamin, 2008, p. 23).

Who Is Talking about Content Operations?

If defining content was not an easy task, providing a definition of ContentOps that satisfies all the disciplines and sectors that have used the term is even more complicated. Part of the problem in defining Content Operations (or ContentOps) is that some see this term primarily as the *operations* of content, including content strategy or content design, while other camps use it to describe the *operationalization* of content, which takes place after content has been designed or a strategy has been developed.

As I write this, there are not many books that provide a definition of Content Operations—although some have been talking about Content Operations for decades but probably not using that term (see Rahel's chapter, “Defining Content Operations,” in this collection). Colleen Jones is a pioneer in this area, and in her book *The Content Advantage (Clout 2.0)*, she argues that “well-executed content operations” should be combined with content vision and strategy to deliver “the right content for the right customer in the right touch-point at the right time” (Jones, 2019, p. xiii). In that same volume, Jones defines Content Operations as follows: “*Content operations* is the behind-the-scenes work of managing content activities as effectively and efficiently as possible. Today, content operations often require a mix of elements related to people, process, and technology” (p. 162).

In 2020, Toby Murdock and Zoë Randolph published the book *Mastering One Voice*, which they describe as “a marketing fable and field guide to content operations.” They provide the following definition for Content Operations: “An internal operating model that moves teams from status to strategy, bringing visibility, collaboration, and confidence to content at scale” (2020, p. 131). In *Leading Content Design*, Rachel McConnell writes that Content Operations “refers to the processes that enable consistent content design” (2022, p. 7). She adds that ContentOps “lays the foundation for effective and efficient work, and frees up team members to think more strategically” (2022, p. 8). And the Scriptorium team, led by our own Sarah O’Keefe, defines ContentOps as “the engine that drives your content lifecycle” (2022, p. 30). Their ContentOps Manifesto—a sister publication of this collection—is a practical, concise, and effective introduction to the topic.

Finding definitions of Content Operations is somewhat easier in the blogosphere and Twitterverse. Some of the contributors to this volume, as noted above, actually started what became their chapters in their personal or corporate blogs. Outside the contributors to this collection, other professionals frequently writing about Content Operations online include Carrie Hane, who has defined the term as follows: “Content operations is the people, processes, and technologies that allow an organization to implement its content strategy to efficiently produce and effectively deliver content” (2019). Hane adds, “Interestingly, most of the definitions came from platforms that enable content

operations for content marketing groups. Yet content is produced and delivered by more than just a marketing team. Like content strategy itself, content operations has to rise out of the silo of one group within an organization to have maximum impact” (2019).

From a software vendor perspective, GatherContent—which advertises its product as a “Content Operations Platform”—frequently tweets about the subject. In its definition, ContentOps “is the combination of people, process and technology that are required to produce, distribute and maintain content in an organisation” (n.d.). Colleen Jones also participates actively in conversations about the term. Her Twitter influence is amplified by the online presences of Content Science and ContentWRX.

It is also worth mentioning what probably could be the oldest definition of Content Operations. Deane Barker, who has been using the term for more than a decade, explained Content Operations in 2016 with the following statement: “Content Operations (CO) is concerned with everything between Content Strategy (CS) and Content Management (CM). Any form of content manipulation and analysis would be managed by a CO process” (Barker, 2016b).

Our Approach

As the subtitle of this volume indicates, the chapters included here are not academic essays or research reports. That is intentional, as the purpose of this book is to present perspectives from industry experts to inspire more advanced conversations and research projects related to ContentOps. Although Rahel has taught for a long time in different academic programs, the only full-time academics in this collection are bookending the essays: I in this introduction, and Jason Swarts in the afterword. Yet through the life cycle of this book (including its very long design stage with Sarah), we always intended to publish it through an academic press to strengthen connections between industry and academia and to incite transdisciplinary dialogue.

The chapters that follow focus on essential considerations for ContentOps implementation, regardless of scale or subject matter involved. First, Rahel defines the term and performs a comprehensive literature review of ContentOps (and other related operational models—or “Ops”). Sarah then focuses on the variables of scalability, velocity, consistency, risk mitigation, and compliance for a given content-related project as indicators to use when making a business case for ContentOps. Sarah establishes some analogies between ContentOps and manufacturing, but Kate Kenyon calls for governance as the key component to ensure that the content machine delivers effective information products for human users. Kate presents a set of checks for standards to keep the content machine efficient but also effective in its processing and publishing work.

Rahel and I then address those readers who are actually content developers. We provide ideas on how to architect content for operating models that demand maximum agility and can scale with ease. Rahel and I introduce a spectrum of robustness for ContentOps, in which different levels enable and allow diverse features and capabilities. At the highest levels of robustness on that spectrum, an attractive feature is content personalization. In order to know the users of a given ContentOps and plan for personalization, a team should conduct research on customer experience, and that is what Kevin P. Nichols addresses in his chapter, arguing that a solid understanding of the customer journey and its stages is necessary to frame its role pertaining to ContentOps. The path to content personalization, once a team has conducted customer experience research, leads to Jeffrey MacIntyre's chapter in this collection. Jeff presents connected experiences as the natural fruit of a strong ContentOps function, which sets up an environment of productive authors creating effective content for personalized delivery that addresses specific users' needs.

Loy Searle explores intersections of localization and ContentOps, and she analyzes their shared publishing systems and emphasis on words that impact customers and their experiences. Patrick Bosek, acknowledging that Content Operations intrinsically relies on systems and tools, gives a platform-agnostic overview of the technology that supports effective ContentOps. Like some other authors in this collection, Patrick makes a living from specific tools and standards, but his recommendations cut across disciplines and silos to deliver guidelines that work regardless of specific software implementations.

In the epilogue, Jonathan McFadden reminds us that ContentOps is, ultimately, about the individuals who consume, use, and interact with the content that we create to accomplish tasks, meet goals, and fulfill obligations. Jon writes about the many faces and needs of those individuals to connect ContentOps to diversity, equity, and inclusion. Lastly, Jason Swarts looks at next steps in the afterward. In particular, he looks at the implications of ContentOps for curriculum in technical communication, with one eye on specialized workflows and knowledge but another on the basic research in the field that provides the foundation for those specialized skills and competencies.

Who Is This Book For?

When Sarah O'Keefe and I started drafting this volume, at one point we envisioned a textbook titled *The Content Professional*. Our original audiences were undergraduate students in technical writing or communication courses (and their instructors) who were interested in current trends in content creation, as well as early-career technical authors or other professionals interested in moving to a career in technical content development.

As the book evolved and incorporated viewpoints from marketing, user experience, content design, and tools development, among others, the audiences

expanded to the following list (most likely incomplete and ambitious at the same time). First, in academia, we aim to reach primarily undergraduate audiences in programs in communication, technical/professional communication and writing, marketing, and computer science, as well as faculty and graduate (or advanced undergraduate) students in the fields of technical communication, computer science, marketing/business information technology, operations management, library science, and any transdisciplinary initiatives related to ContentOps. We also envision an academic audience of community college faculty for applied technical communication courses. In industry, we hope to reach trainers and content professionals and managers working with technical or product content, marketing content, and related technologies. Managers in technical publications teams, technical writers/communicators, information architects, content strategists, content designers, and UX writers will also find this material worthwhile, and we foresee an audience sector represented by information technology professionals interested in contrasting ContentOps to DevOps or ITOps.

I earnestly hope that this collection answers some questions—and inspires a few more—for anyone who is interested in the concept of Content Operations.

References

Abel, S. (2014). Content. In S. Abel & R.A. Bailie (Eds.), *The language of content strategy* (pp. 12–13). Laguna Hills, CA: XML Press.

Andersen, R., & Batova, T. (2015). The current state of component content management: An integrative literature review. *IEEE Transactions on Professional Communication*, 58(3), 247–70.

Atherton, M., & Hane, C. (2018). *Designing connected content: Plan and model digital products for today and tomorrow*. San Francisco, CA: New Riders.

Bailie, R.A. (2014). Content strategy. In S. Abel & R.A. Bailie (Eds.), *The language of content strategy* (pp. 14–15). Laguna Hills, CA: XML Press.

Barker, D. (2016a). *Web content management: Systems, features, and best practices*. Sebastopol, CA: O'Reilly.

Barker, D. (2016b, January 27). The need for content operations. *Dean Barker's Technical Blog*. Retrieved June 17, 2022, from <https://deanebarker.net/tech/blog/need-for-content-operations/>.

Benjamin, W. (2018). *The work of art in the age of its technological reproducibility, and other writings on media*. Cambridge, MA: Harvard University Press.

Bosek, P. (2020, June 11). What is ContentOps? *LinkedIn Pulse*. Retrieved June 17, 2022, from <https://www.linkedin.com/pulse/what-contentops-patrick-bosek/>.

Dush, L. (2015). When writing becomes content. *College Composition and Communication*, 67(2), 173–96.

GatherContent. (n.d.). *What is 'ContentOps' and why should you invest in it?* Retrieved June 17, 2022, from <https://gathercontent.com/content-operations>.

Gollner, J. (2014, April 5). Putting content in its place. *Content and Management*. Retrieved June 17, 2022, from <https://www.gollner.ca/2014/04/putting-content-in-its-place.html>.

Hane, C. (2019, July 25). Content operations—The missing link between strategy and delivery. *Zen—A Tanzen Blog*. Retrieved June 17, 2022, from <https://www.tanzenconsulting.com/blog/2019/7/24/content-operations-the-missing-link-between-strategy-and-delivery>.

Halvorson, K., & Rach, M. (2012). *Content strategy for the web*. 2nd ed. Berkeley, CA: New Riders.

Jones, C. (2019). *The content advantage (Clout 2.0)*. San Francisco, CA: New Riders.

Kimball, M.A. (2016). The golden age of technical communication. *Journal of Technical Writing and Communication*, 47(3), 330–358.

Mathewson, J., & Moran, M. (2016). *Outside-in marketing: Using big data to guide your content marketing*. Boston, MA: IBM Press.

Masycheff, A. (2018, May 22). Keeping technical writing jobs relevant with content automation technologies. *LinkedIn Pulse*. Retrieved June 17, 2022, from <https://www.linkedin.com/pulse/keeping-technical-writing-jobs-relevant-content-alex-masycheff/>.

McConnell, R. (2022). *Leading content design*. New York: A Book Apart.

Murdock, T., & Randolph, Z. (2020). *Mastering one voice: A marketing fable and field guide to content operations*. Boulder, CO: Upland Software.

Pringle, A., & O'Keefe, S. (2009). *Technical writing 101: A real-world guide to planning and writing technical content*. Durham, NC: Scriptorium Press.

Scriptorium. (2022). *Content transformation: An introduction to Enterprise Content Ops and content strategy*. Research Triangle Park, NC: Scriptorium Publishing Services.

CHAPTER I

Defining Content Operations

Rahel Bailie

Content Operations may seem like a relatively new addition to the complement of fields of practice within the content industry. In fact, Content Operations has been around as long as content has been produced as part of an organizational process. It wasn't called Content Operations—or ContentOps, to use the term that's become part of the technology vernacular. Deane Barker, a respected consultant and author of O'Reilly's *Web Content Management* book (Barker, 2016), has found references to Content Operations in one of his blog posts from over a decade ago. Similarly, I have found references to Content Operations in articles and presentations that I created in the early 2000s.

The difference between talking about Content Operations then and now seems to be a change in mental model over a period of time. One reason why the understanding of ContentOps is gaining traction could be concept formation—that is, the process of sorting specific experiences into general rules or classes. Because concepts such as DevOps, DesignOps, ResearchOps, FinanceOps, AIOps, and so on have become familiar, it's not much of a mental stretch to consider what ContentOps might mean. Another reason may be that the need for operational models has grown along with the size of our content corpora. When enough organizations feel the pain of trying to manage a steadily growing body of content, they start to look for ways to operationalize their processes.

How to cite this book chapter:

Bailie, R. 2024. Defining Content Operations. In: Evia, C. (ed.) *Content Operations from Start to Scale: Perspectives from Industry Experts*. Pp. 13–24. Blacksburg: Virginia Tech Publishing. DOI: https://doi.org/10.21061/content_operations_evia_3. License: CC BY

Operating Models

“Content is one of the oldest disciplines, yet it has the least mature operational controls in today’s enterprise.” This is one of two ideas that I often find myself repeating. (The second idea comes at the end of this chapter.) To understand ContentOps, we first need to understand what it means to have an operating model. Let’s look at some definitions.

One definition, from Wikipedia, starts this way: “An **operating model** is both an abstract and visual representation (model) of how an organization delivers value to its customers or beneficiaries as well as how an organization actually runs itself.”¹ The expansion of this definition is telling: the author maintains that “probably the most common use of the operating model tool is to get alignment between managers in different functions or divisions about how they are going to work together for the benefit of the whole.” In other words, an operating model works to define processes, especially across departments (or team, squads, or any other euphemism enterprise teams now use), so that the overall goals of the organization can be achieved.

According to the Strategy& group of PricewaterhouseCoopers (PwC), “An operating model determines behavior, workflow and process design, IT decisions, and investment decisions” (n.d.). The authors emphasize that an operating model must be aligned with an organization’s strategy to be effective. They point out that underperforming operating models may be a detriment to the overall business in three general ways. The first is hoarded or wasted resources; the second is not following through on strategic goals; and the third is a corporate culture that doesn’t enforce accountability. I will return to this statement when looking specifically at the operating models for content.

The title of the EY (Ernst & Young) report *Operating Models: Delivering on Strategy and Optimizing Processes* (Murphy et al., 2016) summarizes its thesis, which is that an operating model is a bridge between an organization’s strategy and operations. The strategy can be understood as *why* an organization is making certain decisions. An operating model determines *what* an organization should be doing to implement the strategy, and operations provides the *how*, so that the implementation can be carried out with consistency.

Operations as the Outcome of a Strategy

How do these definitions play out in the content world? Since Ann Rockley’s seminal book, *Managing Enterprise Content: A Unified Content Strategy*, was first published (2002), the industry has been focused on content strategy as if this concept had just been invented. The reality is that since the 1940s,

¹ Wikipedia, s.v. “Operating model,” https://en.wikipedia.org/wiki/Operating_model, retrieved September 4, 2021.

practitioners from multiple professions have attempted to find operating models for content production. In “An Uneven History of Content Strategy,” Don Day, one of the pioneers of componentized content, comments on the invention of structured markup languages to automate content use and reuse (Bailie, 2020).

The irony is that content strategies were often formulated and implemented by technical writers as a standard expectation of their jobs. Rather than a progression from strategy to operating model and then to operationalizing content processes, the content industry backed its way into articulating strategy. In the early days of content strategy, some writers scoffed at what they saw as an inflated name for something they did regularly alongside the main job of developing content.

Content strategy is the formulation of a plan that codifies a repeatable system governing the management of content throughout the entire content life cycle. In other words, we create a content strategy to find a way to operationalize content processes. The content strategy needs to connect to ContentOps, whether that's planning for more efficiency on the marketing or editorial side or designing a technical content ecosystem to increase delivery efficiencies. Oftentimes, it's both; more efficiency requires an upgrade in tooling. In organizations that have squandered significant time and effort, ContentOps is not a luxury. It's critical for smooth operations and a competitive edge.

Eventually, the recognition that having a content strategy was an important step toward meeting business goals with better content processes led to a separation of the strategy phase, the implementation phase, and the operational phase. Of course, a content strategy is worthless if an organization puts its carefully crafted, elegant strategy into a drawer to gather dust. If there is no implementation phase, then the strategy remains simply an unimplemented plan.

Operational Models across Industry

Before looking specifically at Content Operations, it's worth visiting some of the operational models in the professions with which content professionals work. By comparing operations management models, we can ascertain how their common principles apply to ContentOps.

Some content developers will be familiar with the concept of DevOps, which is a software engineering methodology that unifies software development (Dev) with information technology operations (Ops). According to the Wikipedia definition,² the goal is to shorten systems' development life cycles while delivering features, fixes, and updates frequently and in close alignment with business objectives. The main characteristic of DevOps is to advocate for automation and

² Wikipedia, s.v. “DevOps,” <https://en.wikipedia.org/wiki/DevOps#Definition>, retrieved January 10, 2019.

monitoring at all steps of software development, from integration, testing, and releasing to deployment and infrastructure management.

DesignOps might be another familiar term to content developers working with product designers. The Nielsen Norman Group defines design operations as “the orchestration and optimization of people, processes, and craft in order to amplify design’s value and impact at scale” (Kaplan, 2019). The explanation for the need for design operations is to address challenges such as helping design teams grow and evolve, finding and hiring people with the right skills, creating efficient workflows, and improving the quality and impact of design outputs. Similarly, ResearchOps has developed over a period of a few years in response to the need for standardized processes around planning, carrying out, and applying research. Nielsen Norman’s definition of research operations is “the orchestration and optimization of people, processes, and craft in order to amplify the value and impact of research at scale” (Kaplan, 2020).

In 2014, DocOps made its appearance, aimed specifically to support technical communication in producing product content (see Johnson, 2014). Its goal makes sense, as the function of technical communication often involves managing massive amounts of content with astounding amounts of complexity. Technical communicators have taken on ContentOps out of sheer necessity. In fact, this discipline has been around since the 1990s, reflecting a super-efficient processing of technical documentation. DocOps has evolved to accommodate demands such as multiproduct, multichannel product content.

One of the most pertinent definitions of an operational model is that of DataOps. Wikipedia describes DataOps³ as an automated, process-oriented methodology, used by analytic and data teams, to improve the quality and reduce the cycle time of data analytics. It incorporates the Agile methodology to shorten the cycle time of analytics development in alignment with business goals. Importantly, “DataOps is not tied to a particular technology, architecture, tool, language, or framework.” Tools that support DataOps promote collaboration, orchestration, quality, security, access, and ease of use.

Commonalities of Operating Models

Comparing operating models across the industry reveals a certain number of recurring themes. At a strategic level, the recurring themes could be expressed as a culmination of these goals:

- Activate strategy to drive value. Operating models need to be planned, and planning points to strategy. All planning gets measured against the organization’s strategic vision for what it wants to accomplish.

³ Wikipedia, s.v. “DataOps,” <https://en.wikipedia.org/wiki/DataOps>, retrieved September 13, 2021.

- Improve collaboration across value streams. Essentially, collaborating across value streams takes systems thinking—a way of looking at the big picture so that the most overall value can be delivered. It also prevents uneven production across the organization, such as code being produced faster than it can be tested.
- Automate continuous delivery pipelines. A continuous delivery pipeline is a foundational principle of the Agile operating model. Delivering in sprints allows for a product to be developed over time, with small, regular contributions rather than a long development cycle with a single big delivery at the end. In regulated environments, there may be a final consumer-facing deliverable; internally, sprints allow each piece of the final deliverable to be produced and tested before delivery.
- Enable delivery of information. Efficient operations mean efficient delivery of information, whether that is content delivery, data sharing, or a combination of both. An information bottleneck hobbles a company's ability to operate. Production of information is seen as a value stream, which needs to be part of the same continuous delivery pipeline as the other streams.
- Manage risk. Risk management is an important aspect of compliance in industries where regulation is mandated. Risk management is becoming important in other industries as well, where reputational damage can easily become an expensive proposition.
- Improve innovation. Continual improvement of the operating model requires an organization to be innovative in its production methods. Standing still while competitors find ways to outperform an organization becomes a form of stagnation.

Turning these strategic themes into actionable aspects of an operating model looks quite tactical. Broadly speaking, the operational objectives are these:

- Reduce inefficiencies. Inefficiencies can be found by doing value stream mapping, which ensures that as much waste⁴ as possible is removed from the production process.
- Automate wherever possible. Removing manual steps from a process, even if the steps only add a minute or five, adds up, especially when those steps could be handled in a microsecond by a computer.
- Develop repeatable processes. Clarity around who does what, and when, in a process and who is responsible for each step along the way allows people to proceed with confidence.
- Allow scaling up of outputs. Whatever systems are in place—people, process, and tooling—should support exponential scale.

⁴ There are many articles online about the seven wastes of Lean services, which include delays, duplication, unnecessary motion, unclear communication, inventory, errors that affect customer satisfaction, and lost opportunities.

- Ensure resource availability. This is a key aspect in supporting the strategy and the operating model: planning for sufficient resources to be available to keep the overall system running smoothly.
- Monitor results and create insights. Continual monitoring of the health of the system to maintain an equilibrium of production efficiency is as important as monitoring the health of the product to gather insights and make improvements to it.

What is important to note here is that the emphasis is on efficiency and scale, not on quality. In fact, quality itself is not specifically called out as part of an operating model. If an operating model is not implemented well, there is a risk that it will simply accelerate poor output. However, if the operating model stays true to implementing the strategy, then the assumption is that the operations will accommodate and enable quality output from the implemented processes.

Case Study of Operational Model Improvements

A good way to demonstrate differences in operating models is to show how accounting changed over time. Let's start by looking at how transactions were traditionally recorded and how they are recorded now.

Since 1494, double-entry bookkeeping dictated the way in which accounting transactions were made. A T-account (whose name derived from the T-shape of the ledger) had at least two entries for each transaction, and often more than two entries, as shown in Figure 1.1. At the end of the month, a trial balance was created to check that all the entries were correct. Anywhere from a day to an entire week would be devoted to checking entries and correcting errors. The head bookkeeper or accounting manager would generate some sort of financial statement that would indicate the company's financial position.

This operating model was manual, labor-intensive, and prone to error. Scaling company activity meant adding more accounting clerks (as in Figure 1.2), whose work could be measured in the number of transactions to be recorded. This operating model continued for some 500 years, until accounting systems started to be computerized in the 1950s.

Asset		Liability		Equity	
1,000 500		300 600		100	700

Figure 1.1: T-ledger for double-entry bookkeeping. Credit: Rahel Bailie.

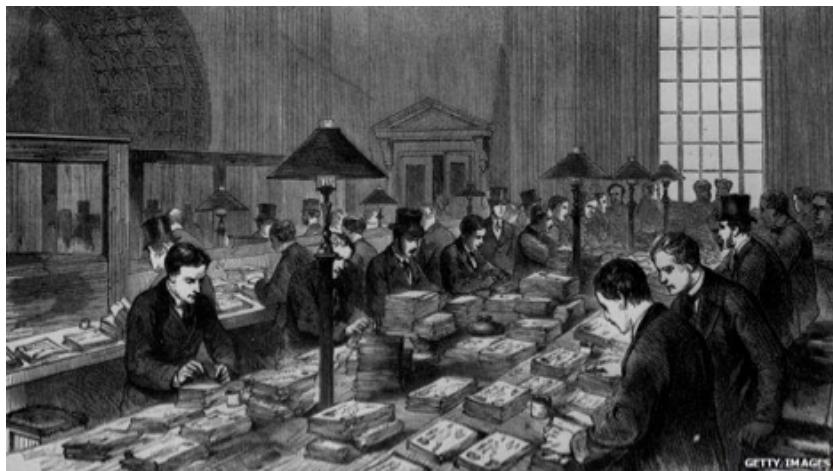


Figure 1.2: Office clerks in Victorian times. Credit: Getty Images/BBC News.

Today's operating model incorporates much greater efficiency. The process uses far fewer people to complete far more transactions. Instead of recording a transaction in multiple places, which then need to be cross-referenced and verified at the end of the month, accounting clerks enter transactions once and let the system handle the rest—if they need to enter each transaction at all. (Many transactions are now transfers of data between systems.) The combination of streamlining, automation, real-time insights, and the number of transactions that can be processed allows companies to scale their capacity efficiently and effectively.

If only organizations valued content assets the same way they value financial assets! Despite Bill Gates's adage that “content is king,” most companies treat content like Cinderella *before* she went to the ball. This collection, *Content Operations from Start to Scale*, aims to elevate content from sweeping the hearth to being part of the royal court.

Operational Models in Content Production

The type of efficiency and scalability seen in the accounting field has been slow to gain momentum in the content field. Practitioners often aren't aware that there are more efficient ways to produce content, so they are not in a position to advocate for a better operating model. Key stakeholders, such as product owners, are generally unaware of what happens during content production, and they frankly don't care as long as content gets into the delivery platform. The technology to support ContentOps has largely not caught up to the level of sophistication seen in other industries. There are pockets of incredible sophistication dominated by

a handful of technology systems. The vast majority of content production, however, has not moved beyond the word processing era of the 1990s. Producers use tools and processes that are meant for casual business, but they are not fit for purpose for the level of complexity now demanded of content production.

Given the framework for operating models across industry, particularly in fields that work with content developers, we can apply the principles to operating models for content. A number of definitions for Content Operations have sprung up (see Mills, 2018), with a few key commonalities:

- The focus is on production processes. Whether ContentOps is called a pipeline, supply chain, or life cycle, the focus is on the mechanics that allow for optimum efficiency. This, in turn, allows content to meet business goals, such as reducing time to market, risk management, ability to personalize, manage multiple languages, and produce at scale.
- An operational model for content relies on having a good content strategy—it's imperative to know the end goal to be able to implement the appropriate operating model to help an organization reach the goal. Cookie-cutter solutions only work for cookie-cutter companies.
- A balanced approach, with a focus on people, processes, and technologies, is needed to ensure successful implementation of an operating model for content. These three factors are often inseparable.

Definition of Content Operations

The definition of Content Operations will surely change over time. At the time of publication, the starting point for a definition that is both tool- and genre-agnostic can be stated as follows: “Content Operations (ContentOps) is the implementation of a strategy that incorporates people, processes, and technologies to optimize content production, so that organizations can leverage their content as business assets while retaining content quality.” This basic definition could be embellished with specific outcomes, such as the ability to deliver personalized content at scale, but that would dilute the essence. The list would get very long and never include all of the possible combinations of strategies, production methods, or business goals. (As an aside, a manifesto from Scriptorium [2022] lays out four basic principles of Content Operations, but the manifesto is an important supplement to the definition, not part of the definition itself.)

Operational Models for Kinetic Content

In reality, most content genres lend themselves to operationalization. How they get operationalized may be quite different, but the drive for efficiency during production through to delivery has propelled the development of solutions.

Table 1.1 shows a few examples of how content within an organization can be produced and delivered more efficiently.

What these kinds of content have in common can be described as kinesis. The content is not published as the end of a supply chain, as happens in print. The content moves during the creation process as it is combined, building-block style, or gets used and reused to form other “shapes” of building blocks. The content moves again during the build process, as content references get resolved and multiple outputs are generated. The content moves again during presentation, with different variants being served up during personalization, omnichannel delivery, and so on.

Table 1.1: A few common content genres, with current approaches and more efficient solutions.

Content genre	Before operationalization	After operationalization
Social media text and images	Write tweets/posts in a document. Copy and paste to multiple social media sites and add images individually. Repeat multiple times a day or week.	Log into social media automation software, create multiple tweets/posts, add images, and schedule for later delivery to multiple platforms.
Product content	Write long-form in multiple documents in multiple folders. Annotate through comments, spreadsheets, and email as approval workflow. Copy-paste into CMS, apps; specialist adds metadata from notes.	Write topic-based content in categorized central repository, add metadata, route for approval workflow using integrated feature, serve via an API into delivery platforms to be pulled by downstream systems.
UI/UX content	Write text in Figma files. Copy-paste into multiple delivery areas without keeping track. Type into how-to instructions.	Create as master references, transclude into wherever needed, for fix-once, transclude everywhere maintenance.
Support content	Write topic-based content in a database, on an as-needed basis, in the customer support silo.	Use existing product content and annotate, as needed, for improving overall support content.
Marketing campaigns	Make arrangements with separate creative agencies. Track progress and transfer files manually.	Track progress of all agencies and their activities in a single dashboard.
Conversation design	Write each line of conversation into cells of a spreadsheet and map the flows in process-flow software; renumber manually with each edit.	As content is written, the flows are created and auto-numbered in real time, with each branch visualized for approval in context.

Any kinetic content needs to be created, built, and delivered in more sophisticated ways than any simple supply chain can handle. Part of the equation is the need for reliability in how other systems can automate the retrieval, aggregation, and delivery of content. For that to happen, the systems need to have a strong indication of what to pick up, when to get it, and where to send it. The other part of the equation, then, is getting that reliability: content developers need to create content in more sophisticated ways, paying attention to the editorial and metadata aspects from the very beginning. And that equation rests on a foundation of technologies that support the content creation and manipulation process (ContentOps) before getting to automation.

Summarizing the Benefits

The commonalities that we see in the content space can act as scaffolding on which we can build an appropriate operating model. To understand how to build a model, we can extend the commonalities to content production.

The common ways in which management frames strategic benefits can be summarized as follows:

- Activate strategy to drive value. Consider the assembly lines of automobile manufacturers or the production floor of a major bakery. Their value comes from solid planning to get the most from their investment. The sophistication didn't happen haphazardly; everything from the level of automation to the placement of equipment has been painstakingly implemented against a strategy. This is no different than developing a content strategy that determines the operating model for efficient production.
- Improve collaboration across value streams. This statement assumes that the organization recognizes content as a value stream, along with code, data, and whatever other streams it considers valuable. Showing an organization's product emptied of content can be a powerful way to demonstrate that content is the primary contact with users—and is thus inherently valuable.
- Automate continuous delivery pipelines. When content needs to be delivered in lockstep with feature development in an Agile environment, there is a need to control the management of content throughout the content life cycle. There is also a need to reuse content whenever possible to contain production time and effort, which is self-evident. The ability to keep content flowing through the delivery pipeline at pace is a key component of keeping the wheels of the production line turning.
- Enable delivery of information. Salim Ismail, lead author of *Exponential Organizations* (Ismail et al., 2014), contends that any organization that wants exponential growth must be able to support scalable delivery of information. At ExO Works, founded by Ismail to teach executives how to grow their organizations and compete in a hyperconnected world, one of the cornerstones for growth is enabling information. When asked by Scott

Abel how to go about enabling information at scale, Ismail's answer was, "That's why you need good content strategists, to figure it out."

- Manage risk. Any organization that is subject to compliance has a vested interest in ensuring that regulators are satisfied that content is being well managed. Aside from this obvious need for an operational model to track content provenance to ensure compliance, organizations outside of regulated industries are seeing the benefits of managing reputational risk. Having a tight audit trail with minimal administrative overhead can be a strong motivator in the risk management area.
- Improve innovation. Management understands that when competitors invest in innovation, standing still becomes the equivalent of stagnation. Innovation comes in many forms, from new ways of communicating with users, to new delivery platforms, to new ways to manage information (content and data) to keep pace with technological developments.

It's all very well to have a strategic vision for improving operational models, but it's also important to be able to articulate the tactical benefits for executing on the strategy.

- Reduce inefficiencies. Once an organization has measured the inefficiencies, it can calculate the benefits, whether that is time and resources saved in content production or other measures related to the overall delivery of the product or service.
- Automate wherever possible. "Stop using people as slow computers" is the second idea that I alluded to at the beginning of the chapter—one of the ideas I find myself repeating. The benefits of automation can range from the obvious, such as time savings, to ancillary benefits, such as error reduction.
- Develop repeatable systems. The benefit of having a strong governance system propagates throughout the organization.
- Allow scaling up of outputs. The benefits of an organization being able to add more products, more product lines, more output channels, more markets, more languages, and more audiences are invaluable.
- Ensure resource availability. Aside from the obvious financial benefits, knowing that an organization can plan for surges in resource demand provides peace of mind and the ability to operate with confidence.
- Monitor results and create insights. This facilitates a continuous improvement cycle for content performance.

References

Bailie, R.A. (2020, September 2). An uneven history of content strategy. *Medium*. Retrieved June 17, 2022, from <https://rahelab.medium.com/an-uneven-history-of-content-strategy-d514cf7eee5>.

Barker, D. (2016). *Web content management: Systems, features, and best practices*. Sebastopol, CA: O'Reilly.

Ismail, S., Malone, M.S., & Van Geest, Y. (2014). *Exponential organizations: Why new organizations are ten times better, faster, and cheaper than yours (and what to do about it)*. New York: Diversion Books.

Johnson, T. (2014, October 10). DocOps: Interview with Jim Turcotte. *I'd rather be writing*. Retrieved September 13, 2021, from <https://idratherbewriting.com/2014/10/21/docops-interview-with-jim-turcotte/>.

Kaplan, K. (2019, July 21). *DesignOps 101*. Nielsen Norman Group. Retrieved June 17, 2022, from <https://www.nngroup.com/articles/design-operations-101/>.

Kaplan, K. (2020, August 16). *ResearchOps 101*. Nielsen Norman Group. Retrieved June 17, 2022, from <https://www.nngroup.com/articles/research-ops-101/>.

Mills, R. (2018, November 26). *Understanding content operations—Definitions and resources*. Gather Content. Retrieved September 13, 2021, from <https://gathercontent.com/blog/understanding-contentops-definitions-and-resources>.

Murphy, A., Kirwin, J., & Abdul Razak, K. (2016). *Operating models: Delivering on strategy and optimizing processes*. Calgary: EY.

Rockley, A. (2002). *Managing enterprise content: A unified content strategy*. San Francisco, CA: New Riders.

Scriptorium. (2022). *Content transformation: An introduction to enterprise content ops and content strategy*. Research Triangle Park, NC: Scriptorium Publishing Services.

Strategy&. (n.d.). *Our solutions—Operating model*. Retrieved June 17, 2022, from <https://www.strategyand.pwc.com/gx/en/unique-solutions/capabilities-driven-strategy/operating.html>.

CHAPTER 2

The Business Case for Content Operations

Sarah O'Keefe

We have an ingrained mental model of writers as introverted hermits, toiling away in solitude. Eventually, they produce manuscripts, which are fed into a publishing pipeline for editing and production. This model might hold for some fiction writers, but content production looks very different for marketing and technical efforts.

Today's corporate content requires close collaboration across multiple specialties, style guides, standardized processes, governance, and industrial-grade tools. Creating content for a large organization resembles a manufacturing process rather than our traditional model of heroic solo writers.

There is an additional complication. Most content is not just written, processed, and delivered once; rather, it undergoes edits, updates, and corrections over time. Although you may package and deliver information, the process doesn't end there. Content production is a life cycle in which information is constantly evolving.

We can borrow further from manufacturing and think of ContentOps as an assembly line, which lets an organization optimize each component of the content development process. Just remember that our content process, unlike an actual assembly line, can loop back on itself for content updates. The idea of a "content factory" is in stark contrast to the image of a solitary writer, and it can provoke resistance or outright hostility. Typically, it's easier for more technical content creators—technical writers, UX writers, and API documentation writers—to think in manufacturing terms than it is for more creative writers in marketing roles.

How to cite this book chapter:

O'Keefe, S. 2024. The Business Case for Content Operations. In: Evia, C. (ed.) *Content Operations from Start to Scale: Perspectives from Industry Experts*. Pp. 25–32. Blacksburg: Virginia Tech Publishing. DOI: https://doi.org/10.21061/content_operations_evia_4. License: CC BY

Executives view ContentOps through a different lens: they demand a business justification for any investment. (For a more detailed discussion of the business drivers, refer to “Summarizing the Benefits” in the “Defining ContentOps” chapter included in this collection.) These are the most common justifications:

- Scalability
- Velocity
- Consistency
- Risk management and compliance

Scalability

In most small businesses, content development is inefficient and fragmented. As long as the content volumes are relatively low (and all in one language), this inefficiency is reasonable.¹

However, as a business grows, content demands multiply. In particular, once a business starts expanding product lines and globalizing, it faces the following content challenges:

- The problem of controlling content without creating multiple copies of information for related documents. (For instance, a company might split a product into “regular” and “premium” versions. The “premium” product has a few additional features, but the overlapping features are the same. For this, the organization needs a way to manage the overlapping content without making copies.)
- The challenge of delivering information into multiple channels—web, print, email, social media, and so on—with consistent messages and terminology.
- The challenge of sharing information across multiple content types—product reference, knowledge base, training, marketing, and so on—while avoiding duplicate or contradictory information.
- The rising demand for locale-specific content, including new languages and potentially new regulatory requirements.
- The need to identify and target specific audiences with certain content.

Thus, a single piece of content might live in multiple product versions, channels, content types, and languages. At this point, the price of fragmented content rises to an unacceptable level. When every piece of content goes to the web first,

¹ Scriptorium uses a guideline of around \$250M in revenue for US organizations. That is the point at which organizations start to prioritize global operations and therefore scalability. The cutoff tends to be lower for European organizations (which prioritize localization earlier in their business life cycle).

is used in several other delivery channels, and is translated into dozens of languages, any friction in the content process gets multiplied for each channel and each language. Five minutes of manually moving content from point A to point B doesn't sound like much, until you have six channels (30 minutes) and 20 languages (600 minutes, assuming five minutes per language per channel). Suddenly, you've spent hours just moving files around. Industry conversations mention that technical writers spend nearly half their time on "document maintenance" tasks.

To build out scalable Content Operations, an organization will need to invest in the following:

- A centralized repository for content.
- A reuse strategy to identify and manage reusable components across channels and content types.
- A conditional strategy to identify and manage variant content (where most of the content is reusable, but one small chunk is unique to a specific deliverable or audience).
- An omnichannel delivery pipeline.
- Audience profiling.
- A global content strategy to address locale-specific content requirements and localization/translation work.

The payoff for this investment is a content pipeline that lets the organization maximize the value of each piece of content.

Velocity

Content velocity affects an organization's ability to speed up time to market, and ContentOps provides a way to improve content velocity. A basic content life cycle looks like this:

Authoring > editing > review/approval > localization > rendering > delivery

In a paper-based life cycle, most advances were confined to physical production, such as faster printing presses. But in a digital content life cycle, an organization can automate and optimize all stages of the content life cycle.

Authoring and Editing

Organizations can speed up authoring either by making content creators more efficient or by reducing the amount of content that needs to be created.

To increase efficiency in the work of content creators, you can provide authoring frameworks, efficient authoring tools, and authoring support (through software and processes). For example, an organization might have a tool that automatically identifies overly complex sentences and offers recommendations to simplify them. Structured content templates can guide authors as they create content to ensure that they include all of the needed components in a particular document. If a magazine article requires a summary at the beginning and an author profile at the end, a structured content tool can prompt authors to include that information.

A reuse strategy, which makes content more scalable, also reduces the amount of content that needs to be written and thereby improves velocity. To increase content reuse, it's critical to provide authors with a way to locate existing content and identify good candidates for reuse. Authors shift from prioritizing writing new content to identifying new ways to mix and match existing content. Most organizations can expect at least 20% reuse in their product content; that number can rise to as high as 80% for certain industries (the semiconductor sector is a good example) that have huge content volumes and lots of overlap among product lines.

Reuse also improves outcomes downstream in the content life cycle—more reuse means less information to edit, review, approve, translate, render, and deliver.

At a bare minimum, editing content by passing around files and using some sort of change tracking is a huge velocity win over paper-based comments. To increase editing velocity, organizations can augment human editors with software for structure and terminology. Another approach is to step away from the author/editor framework and instead create shared documents for collaborative authoring. If a group of two or three authors work together in a shared file (Google Docs is a great example of this), they can create a collaborative document instead of each working on their own personal filesets. A truly collaborative writing approach blurs the distinction between authoring and editing.

Maintenance

Once content is published, it needs to be maintained. Typically, that means correcting any errors and making updates as things change. A solid ContentOps workflow means that you can update a piece of content in a single location and have the change flow to every place that uses that information. If content is reused via copy and paste, then a single content change needs to be made in multiple locations. Those problems multiply across languages and content variants.

Another opportunity in maintenance is to examine how changes and corrections are captured and managed. For example, how are user comments handled? Are they ignored, or is there a process to capture them, validate the information, and then ensure that the underlying content is updated? After the correction is made and published, what should happen to the comment?

Review and Approval

The review and approval process is a common cause of friction in the content life cycle. The problem often lies with limited authority. If a single person is responsible for approving content, that person's availability determines how quickly content moves through the approval process.

The single point of failure problem can be addressed by increasing the number of people who have approval authority or by identifying a backup approver when the primary approver isn't available. Once an organization has clarified the approval assignments, it should consider a review and approval workflow, which may live inside its content management system. This software lets the company set up assignments and notifications, so that when an author completes a piece of content, the content is automatically routed to reviewers. Reviews could be serial (reviewer A, then reviewer B, then approver C) or parallel (reviewers A, B, and C all review at the same time, and when their issues are resolved, the content moves into an approved state).

Review and approval workflows vary widely across industries and organizations. In some places, authors approve and publish their own content. In other organizations, extensive review cycles are the norm. Regulated industries typically have compliance requirements that drive their review process. Review stakeholders may also include legal teams or quality assurance.

Rendering

Velocity in rendering requires formatting automation. Content is stored with tags or labels that indicate meaning (like “heading 1,” “button label,” or “warning”), and then the appropriate formatting is applied as the information is rendered for PDF, HTML, or other formats. A multichannel delivery pipeline requires the organization to think about rendering across many channels and ensure that the content has all of the labels needed to create every format.

For maximum velocity, a content team needs to ensure that all rendering is automated. Furthermore, it should build in localization support for all target languages. Manual formatting is doable in small ContentOps, but it will become a problem as the organization scales.

Delivery

Delivery is perhaps the phase that has been most transformed by the shift from paper to digital workflows. Although modern ContentOps workflows have added new tools and technologies everywhere, authoring and editing is still recognizably the same process on paper as in a digital workflow. But delivering

paper documents requires manufacturing (to create physical books) and logistics for actual physical delivery, as opposed to putting content on a website to make it available worldwide in a split second.

So even without formal ContentOps, digital delivery is faster than physical delivery. The content team does end up with complications because the number of channels that they need to deliver to has increased. ContentOps for delivery requires thinking carefully about content governance—how soon after approval should content be posted? Is there a delivery schedule? Do you use content delivery networks or other intermediaries to manage the load?

Another way to look at delivery is to use a *pull* rather than a *push* model. Instead of finalizing content and then pushing it to publication channels, an organization can have content clients. The content client requests information from the organization's content repository (or an intermediate layer) and renders the content that's delivered to the client.

Digital delivery should be instantaneous in any digital workflow, so once a team gets to this point, it doesn't have to worry too much about velocity.

Consistency

Improving consistency of content provides another justification for investing in ContentOps. The technology and processes in a mature Content Operations environment make it easier to achieve the following:

- Control terminology across all languages.
- Ensure that the look and feel of content matches; for example, if the rule is to italicize glossary terms in technical content, then all glossary terms are italicized throughout the content set.
- Ensure that content used in multiple places is the same throughout the corpus.

Content consistency helps build user trust and makes it easier for users to understand information. In high-stakes content, such as that related to medical devices or industrial equipment, content consistency helps ensure the safety of the people using the products. Ensuring that all warnings are highlighted consistently and follow industry standards helps people avoid injuries due to incorrect product use. (It may also reduce the manufacturer's legal liability if an injury does unfortunately occur.)

In addition to safety issues, consistency helps with brand identity and customer trust in the following areas:

- The use of consistent terminology across all channels and content types makes the customer feel more comfortable. Customers build confidence as they learn an organization's terminology, instead of stumbling when multiple terms are used for the same concept.

- A consistent look and feel assures the customer that the content is trustworthy. When clients notice design variances, they may wonder why they occurred. Does inconsistent design mean that the content is not fully vetted?
- Consistent voice and tone help support the brand identity and messaging.
- Consistent design patterns (for example, warnings always boxed in red) mean that customers get familiar with a team's design and know how to navigate the content.
- Consistent content writing can be reused across multiple channels and content types, which reduces the overall cost of ownership for the content.

The business justifications for consistency run the gamut from “stay in compliance with regulators” to “build trust in our brand.” Each organization will value consistency based on different considerations.

Risk Management and Compliance

I've mentioned risk management and compliance as a factor in several of the other business justifications, but I think it's worth addressing separately. If an organization has compliance requirements, ContentOps can formalize the content life cycle and reduce the risk of compliance errors.

Providing the wrong content or omitting a required content component in a regulated environment can lead to delays in product approvals, fines, or worse. Establishing rigorous ContentOps that prevent these errors is well worth the cost because the risk is so high. Even without compliance requirements, better ContentOps is a risk-mitigation strategy. If an organization has good control over its content, consistent formatting, and appropriate reuse, it reduces the risk of content errors.

Publishing content introduces some risk for any organization, but it is especially important for regulated organizations to get their content right. For example:

- Submitting incorrect information to a governmental body could result in sanctions.
- Having policies and procedures that do not accurately reflect how a medical device manufacturer operates could result in the factory being shut down.
- Incorrect operational instructions could result in product users being injured or killed.

Risk mitigation is more important in some industries (like industrial equipment) than others (such as video games), so each content team should consider the risk profile for its products.

Building Your Business Case

A scrappy startup with a couple hundred pages of content in three languages needs a different solution than a global medical device manufacturer, and the investment should be commensurate with the expected returns. So as you build out ContentOps, assess your organization's requirements for scalability, velocity, consistency, risk mitigation, and compliance—and build accordingly.

CHAPTER 3

Governance for Content Operations

Kate Kenyon

Digital governance is a framework for establishing accountability, roles, and decision-making authority for an organization's digital presence—which means its websites, mobile sites, social channels, and any other Internet and Web-enabled products and services.

—Welchman, 2015, p. 11

When you think about Content Operations, your mind probably doesn't jump immediately to governance. You're much more likely to think about process, or even the outcomes of ContentOps: making things go more smoothly, more efficiently, or even just faster.

But as Rahel Bailie noted in her chapter in this collection, operations has an “emphasis … on efficiency and scale, not on quality. In fact, quality itself is not specifically called out as part of an operating model. If an operating model is not implemented well, there is a risk that it will simply accelerate poor output.” So how does a content professional stop the machine from turning out bad content—only faster?

The Role of Governance in Content Operations

I would argue that part of a ContentOps model being “implemented well” is having enough governance to know that what an organization is putting

How to cite this book chapter:

Kenyon, K. 2024. Governance for Content Operations. In: Evia, C. (ed.) *Content Operations from Start to Scale: Perspectives from Industry Experts*. Pp. 33–42. Blacksburg: Virginia Tech Publishing. DOI: https://doi.org/10.21061/content_operations_evia_5. License: CC BY

out meets the organization's definition of successful content. Governance underpins an Ops model: it decides the rules and standards for the content. Operating models are how an organization implements those rules and checks that those standards are being followed.

Everyone in a company has a mental model of what good content looks like and how it behaves. In regulated industries, such as banking and pharmaceuticals, there are also external codes and models of how content should behave—coupled with large, painful fines for breaching these standards.

Without clear governance and standards, there's no strong consensus on what that "good content" is, and as time goes on, this leads to expectation gaps between what is expected vs. what is being delivered. An expectation gap sounds like "I just thought it would be easier to do" or "I thought it wouldn't take this long," or even "we followed the process right, but I just don't think it's very good."

These gaps are enough to upend any operational changes that an organization is making. No matter how fast you can get the machine to run, people will pull the factory brake if they don't like the end results.

So: governance as an underpinning for an operations model. Where to start?

I've structured this chapter as a set of questions to help you consider your organization's needs. Not all will apply, so skip anything that doesn't work for you. Governance needs change over time, so if you do skip things now, remember to come back and review!

What Do You Need to Govern?

This is the starting point: what do you even define as "content"? We all think we know, but is there an agreed-on definition across the organization of what content includes? It may just be "all the words, images and documents associated with our products," or it could be broader than that and include app microcopy, customer service content, and social media posts. Consider also content that you lease or buy from third parties, such as newsfeeds.

Defining what you mean by "content" is the beginning of deciding the standards that you will need. As with all parts of governance, it's not a once-and-done thing—if your scope of content changes, you may need to revisit your standards.

What Does "Good Enough" Look Like?

The first thing here is a little bit of organizational politics. Everyone has to largely agree on what "good enough" looks like, and everyone has to agree that this standard is worth working hard to meet. Without those two agreements, it's very difficult to define effective ContentOps and get people to stick to it. The

organization will soon devolve to just-do-it processes where the person who shouts the loudest gets their stuff done by whatever means necessary.

Industries with significant fines/regulations for misleading content (e.g., pharmaceuticals, banking) often have a head start in these areas because the standards are usually compulsory, and the costs of not meeting them are punitive. But even outside these industries, there are standards that companies can agree to for content: accessible content, accurate content, portable content, timely content, and content in the right languages, for starters, and that's before an organization gets to the fine details of brand voice, tone, and suitable imagery.

To give you a rough yardstick: within banking, I would consider the following inputs into content to be “good enough.”

1. Meeting federal banking and financial regulations.
2. Meeting internal standards for risk, legality, and fraud.
3. Meeting audit and compliance standards.
4. Meeting Web Content Accessibility Guidelines (WCAG).
5. Allowing open banking interoperability.
6. Establishing a consistent brand voice and tone.

These standards should be clearly mapped to existing examples and documentation; having content that is “on brand,” for example, is not enough. The implementation of standards should go further and adhere to the organization’s brand book and guidelines. If you don’t have that level of specificity, you will struggle to make the rules clear enough to follow.

Who Gets to Decide Your Content Definitions and Standards?

This question will inevitably pop up as an organization starts defining its content standards, so it’s worth tackling it early. Everyone has an opinion, but whose do you need to set a standard? Does everyone get a say? I would suggest that *not* everyone gets a say. Whoever is involved should have skills and expertise to offer in defining the content standard; they should not be consulted simply because they’re senior and would like their opinions included. The standards will certainly need senior stakeholder support, but that is not a necessary step.

Content has functionality—it should *do something*. Some content is very valuable to your business, and some less so. Every time you publish something that you know isn’t very good (but someone insisted on its being there), you are diluting the chances of your customers finding the *valuable* content. My recommendation is to start by setting a standard around what gets created, so you only have to manage content that is worth having. Choose people who know what “good” looks like and who know what is possible within the organization *right now*. Standards should be achievable outcomes, not aspirational daydreams.

How Much Governance Do You Need to Make Your Standards Stick? And Where Do You Apply It?

Once you have a scope of content, a definition of “good enough,” and the right people to govern it, you get to the thorny part—what processes and tools are you going to use to make sure that these standards are met?

This is where theory meets practice. We all mean to be good people who do the right thing. And yet, if time is short and people are tired, overworked, underskilled, or just don’t see the point, then standards will be the first thing to go. So *where* do you apply the most scrutiny to your content? At the beginning of the process, in pre-production? During production? Or post-production, once it’s live and visible to customers?

It’s likely to be a mix of all three, but it’s not likely to be an equal split. The areas where your content needs most governance will depend on the role that content plays in your organization’s commercial success, the volume and velocity of content production, and what happens if the content isn’t *right*.

For example, the content governance needs for a fast-fashion e-commerce clothing company will be very different from those of a bank offering financial insights about asset and wealth management. Accurate imagery for clothing is key to sales and is also relatively easy to govern and check during production. Conversely, images have limited impact on the success of the bank’s content, but if the wording isn’t accurate, there are significant legal risks—so let’s get that content approved by the right people *well before* production begins.

Some examples of governance standards within processes at each point are as follows.

- Pre-production
 - Suitability: do you even need this content? To determine that, consider that all content should have the following:
 - A clear purpose
 - An outcome for the business and the customer (and no, “awareness” isn’t an outcome)
 - An end date
 - A proven need to exist
 - Ownership: who is going to own this content throughout its life cycle? Is it the person who commissions the content, or the person/team that delivers it? Without a clear view of who will be responsible, this task tends to fall to the “last person (or team) who touched it,” and that’s rarely the person who actually commissioned the content in the first place.
 - Review date: how long is this content for? Content is rarely as evergreen as we think. It needs an agreed-on shelf life, with regular review cycles to sniff it for freshness, relevance, and accuracy.

In my practice, I have found content *eight years* out of date on banking websites. I have found terms and conditions with vital links that no longer work, possibly invalidating the terms themselves. This is what happens without content ownership, accountability, and management—which stems from a lack of pre-production governance.

My personal governance preference is to make those who commission any content the owners of that content, responsible for final say about when it goes into production and responsible for monitoring it once it's live. They are also the “owners” and will be asked to review that content every week, month, or year to ensure that it's still fit for purpose.

This approach creates a clear matrix of responsibility between those commissioning content and those creating, building, and maintaining it.

- Production
 - Consider your creation workflows: what's the standard for your content? If you have already met the pre-production standards of knowing why you're making this content and who it's for, the briefing process should be pretty clear. Woolly briefs reflect woolly thinking, so if you're seeing this at the point of production, it's a good sign that standards aren't being defined or met earlier on.
 - Approvals: during production, your ContentOps workflow needs regular checkpoints, and each checkpoint needs a clear governance remit. What's being checked, why, and by whom? Table 3.1 gives an example of some possible checkpoints.
 - When setting checkpoints, consider how you will get the content to the correct people. Can you give them access to your systems, or is that rather risky? Consider how you can set up a shared space to gather and record feedback, particularly if you're working with regulated content.

Table 3.1: Possible checkpoints for governance review in a ContentOps workflow.

Type	Scope	Reviewers
Web page	Layout check	<ul style="list-style-type: none"> • Content owner • Brand team • Design team • Front-end developers
Email	Pre-send check	<ul style="list-style-type: none"> • Content owner • Marketing team • Accessibility specialist • Data team
Legal document	Final check	<ul style="list-style-type: none"> • Content owner • Legal reviewer • Product owner

- If you need to ask a lot of people for feedback (bad luck!), then consider the impact of this on your production times. How long can people review the material before they need to respond? Is this a person's sole job, or is this a request that you're making of someone who also has other things to do? Set a deadline for people—and be prepared to chase them.
- Post-production
 - Post-production governance is more about quality control than anything else. It involves sampling content to see whether it meets specific criteria—if something is there or not, or whether it is correct or incorrect. As this is a rather binary approach, you can consider automating post-production checks, and there are a number of products on the market that offer web- and document-checking capabilities. These products will check for problems in the following areas (and others).
 - Accessibility: does your web content correctly use H1 and H2 titles for level-1 and level-2 headings, respectively? Does it meet the WCAG guidelines for accessibility?
 - Grammar: is your content grammatically accurate (in all languages)? Does it use a consistent tone and voice?
 - Performance: are your web pages loading quickly and accurately? Are you meeting search engine optimization (SEO) standards?
 - Metadata: does your content accurately and adequately use one of the many published standards for metadata?
 - Images and video quality and sizing: are these correct in your content? Do they have alternative text tags and descriptions?

This is a checklist approach to hygiene, but it's a useful monitor for your production standards: if those are slipping, it will show up in these metrics. It's also particularly useful if you have a large volume of content to govern and only a few people can make it part of their daily responsibilities.

For additional finesse, I would personally recommend periodic audits of existing content for quality, suitability, and tone. These audits are done manually but can give a wider view of the overall body of content rather than individual pieces.

The focus of post-production audits needs to be carefully identified. (What are you reviewing? What are you benchmarking against?) But audits can keep content production values high. I have done this kind of audit over a body of content (e.g., customer emails) to show where there's too much variance in writing styles. Such differences produce a disjointed customer experience. This type of audit is particularly worth doing when an organization has made significant changes to its products or its designs. Sometimes content becomes stale simply as the company evolves.

Post-production governance should also factor in how and when content is released “into the wild.” Treating content like code means that the material will

be checked thoroughly for errors before it is released—but on the flip side, it also means a loss of speed if your release process isn't super-slick. Relying on post-production release mechanisms to catch issues will directly influence how your technology stack is configured, so you'll need a tech partner to help you work through these issues in order to understand what will be best for your organization.

Do You Need Some Nuance in Your Content Governance? Is All Your Content the Same?

In the rush to define a process, differences among types of content are often overlooked. How might those differences affect content governance? It's fine to start by saying "these are the rules for our content," but if it takes three weeks to get a tweet signed off (true story!), then you will fairly quickly need to decide whether all content should be governed in the same way.

When looking at where content governance needs to be altered, here are some key areas for consideration.

- **Timeliness:** some content platforms move more quickly than others. Social media, in particular, relies on being timely, relevant, and responsive. Options for social media governance include managing risks through limiting what subjects can be talked about, having a quicker approvals process, or having a much lighter governance policy that accepts potential risks in favor of the brand value that comes with responsiveness.
- **Urgency:** having a governance plan in place for how to publish content in emergencies, such as natural disasters, major company issues, and the like, is really important, even if you never have to invoke it. Know who your people are for rapid communication, and make sure they know that they need to respond quickly.
- **Risk:** this could be financial, regulatory, or reputational risk. Some content is particularly contentious. One way of viewing content governance as it relates to risk is through the lens of "how badly could this go?" Perhaps you can apply less stringent review processes to low-risk content. Organizations will often have to feel their way along as content governance develops. At first it all feels risky, but as people become more fluent and familiar with standards, spaces for change will open up.

Global vs Local: Do Your Standards Apply Everywhere?

This question will come up even if your content is only in one language! As language reflects the culture and expectations of the people who speak it, what

is “good enough” in one region or country may not cut it elsewhere, even if the words are readily understood. Voice and tone guides need to be reviewed to see whether they will scale to meet the expectations of a new audience. Legal and operational differences may mean that content is not even suitable for reuse but must be reimagined for a different region or country. Even emojis need consideration—an innocuous thumbs-up, for example, is seen as offensive by many in Greece and in Middle Eastern countries (Rawlings, 2018).

If you intend to share and reuse a lot of content across multiple countries, you may need to review the balance of your overall governance model entirely. Unless everyone is polylingual, centralized governance across multiple countries can involve content being translated (or back-translated) in order to be reviewed.

If a centralized model of governance becomes too cumbersome and nobody can get content out in a timely fashion, it’s time to rethink. In these cases, a federated model of governance is often more effective. The federated model involves a central team for policy and standards; local teams, with local knowledge, are empowered to decide how policy and standards are implemented. This model requires trust and collaboration but is generally more effective in business. I’ve seen it used in large marketing departments of corporations as well as within product design systems.

What’s Your Plan for the Governance Life Cycle?

As I mentioned in the introduction to this chapter, governance is an ongoing process rather than one-and-done. If governance is to stick and be adhered to, it will take more than an initial enthusiasm for “good content” to make that happen. Successful application of governance standards requires those standards to be baked into the processes of content creation and the technology that supports them. And to keep your organization honest, you will also need to consider how to monitor whether your agreed-on standards are being met.

The threat of external review and censure ensures that those organizations in regulated industries recognize the need to self-police. But if your company is not externally regulated, who and what could you use to monitor governance adherence? The most common answer is a periodic audit of governance that looks at a sample of n pieces of content to see whether the agreed-on standards were followed—and if not, why not.

These periodic audits are often sprung upon teams to avoid the temptation to cover up problems, but checking for governance adherence should not be used to punish those who don’t follow the rules. Instead, introduce audits as a temperature check that will tell the organization whether the governance is healthy or needs an update.

Another approach to governance review is to schedule it (every three or six months works well in most companies) and then to encourage people to bring

their process complaints to be aired, reviewed, and incorporated into newer standards. I have seen this approach succeed, but only when individuals have a certain degree of trust in their companies and feel that they have the ability to speak up. If proposing new governance is seen as “rocking the boat,” or if governance is understood as set in stone, this approach won’t uncover issues that need to be addressed.¹

What’s Your Plan if Things Go Wrong?

Not to end this chapter on a down note, but part of governance involves planning for when governance fails. It’s planning for when confidential content ends up being shared publicly, when a disgruntled employee leaks documents, when content is published too early or too late, when it accidentally breaks copyright, or when it unintentionally causes offense. You can look to newspapers and social media for examples of your competition experiencing such failures!

Running disaster exercises will allow you to plan, somewhat, for resilience in times of crisis. Be sure that you have a map and processes for what should happen if something goes wrong with your content. Who gets notified? How? And how long will you give the organization to fix the problem?

A Conclusion

Within all ContentOps models (or operating models for content production), the inevitable question that comes up from stakeholders is “yes, but how long until I can see my stuff live?” Content governance sets the standards that help technologists and strategists answer that question. Knowing what is “good enough,” and how and where that standard will be applied, provides the basis for getting your business that magical service-level agreement.

Without content governance, ContentOps becomes purely about process speed and efficiency, and it disregards quality. If a company decides to pare back process without reference to standards, it will inevitably get back into a position where gaps arise in expectations for content. Ensure that your content standards are defined, realistic, regularly reviewed, well documented, supported through technology, and seen within your organization as worth delivering. The standards will provide the *why* for much of your ContentOps as well as the definition of success for your content. “Done” is not the finish line for content—“good enough” is the mark to aim for.

¹ A third approach could be to use automated content optimization software, which Sarah O’Keefe mentions in her chapter, “The Business Case for Content Operations,” in this collection.

References

Rawlings, A. (2018, December 11). Why emoji mean different things in different cultures. *BBC Future*. Retrieved June 17, 2022, from <https://www.bbc.com/future/article/20181211-why-emoji-mean-different-things-in-different-cultures>.

Welchman, L. (2015). *Managing chaos: Digital governance by design*. Brooklyn, NY: Rosenfeld Media.

CHAPTER 4

Operationalizing Content Creation

Rahel Bailie and Carlos Evia

In this volume, we have now defined Content Operations, established the context for effective ContentOps, introduced the business benefits that can be realized by having the right operating model, and discussed the importance of the governance aspects of content-related projects. By now, you are probably wondering about the actual content development—from writing to tagging to managing and eventually delivering content.

This chapter addresses the actual creation of content. However, it is not a detailed guide on how to write or develop content. Instead, it provides guidance on how to architect content for operating models that demand increasing levels of agility in exchange for the ability to scale with ease. We refer to this specific combination of demand and exchange as the level of robustness of a content practice. As a concept, robustness can mean different things depending on the discipline or context in which it is being mentioned. For example, in structural engineering, robustness is a desirable property in structures and systems. It implies “tolerance to damage from extreme loads or accidental loads,” but it is also “applicable to other adverse effects such as sensitivity to human error and deterioration” (Baker et al., 2008, p. 253). Other descriptions of robustness present it as “a design principle of natural, engineering, or social systems that have been designed or selected for stability”—a definition attributed to the Santa Fe Institute (Baker et al., 2008, p. 254)—and “the ability of a system to continue to operate correctly across a wide range of operational conditions, and to fail gracefully outside of that range” (Gribble, 2001, p. 21).

How to cite this book chapter:

Bailie, R. and Evia, C. 2024. Operationalizing Content Creation. In: Evia, C. (ed.) *Content Operations from Start to Scale: Perspectives from Industry Experts*. Pp. 43–74. Blacksburg: Virginia Tech Publishing. DOI: https://doi.org/10.21061/content_operations_evia_6. License: CC BY

In the context of ContentOps, robustness inherits components of all of these definitions: a robust content workflow needs to be tolerant of adverse effects, stable, and able to operate across a wide range of conditions or degrade (fail) gracefully.

We are aware that not all ContentOps require the same benefits and affordances, and in this chapter we introduce several considerations based on the level of robustness set for a specific content-related project. These considerations can be summarized as the degree to which content production facilitates ContentOps (which we describe as the level of kinesis for a piece of content), the adoption and implementation of editorial conventions and technical standards to produce intelligent content, and the establishment of a ContentOps North Star: delivering content as a service (instead of publishing as a deliverable or set of deliverables).

Producing Content to Facilitate Content Operations

For content to be able to work within an operating model that has any level of robustness (or any degree of complexity), the way the content is produced and managed needs to be carefully planned and executed. Content needs to be *kinetic*—to move from state to state along the content supply chain. Content may exist across multiple repositories, with multiple variants, in multiple languages, and end up in multiple outputs, in multiple formats or media. Even formerly simple operating models—e.g., publishing blog posts—can now involve the inclusion of content blocks such as feature lists or quotes from external sources and pulling out social media posts. In operating models for products, digital or analog, there is a great need for content that has some sort of kinesis. We propose the following principles for creating kinetic content that facilitates ContentOps:

- Focus on information enablement.
- Automate continuous delivery pipelines.
- Improve collaboration across value streams.
- Make semantics the foundation.
- Consider technical capabilities.

Focus on Information Enablement

Information enablement means having full command of information and being able to leverage computation power that can scale infinitely.

Information enablement is evidenced in several ways. Salim Ismail (2020) lays out three forms of information enablement—each of which can be applied to technical communication products and processes:

- Turning traditional functions from analog to digital. Rather than distributing “documents” with a tl;dr (too long; didn’t read) amount of information, enable information to be surfaced to the right audiences, at the right time in the user journey, in the right language, in the right format, and so on. An example would be to show explanations of automobile functions in an augmented reality app rather than in a printed manual.
- Harnessing information at scale. Organizations need to utilize information and data produced as needed and at scale. In technical communication, this means creating content in small, discrete units so that it can be auto-aggregated into new information assets, with content variables that can be auto-populated as part of the delivery process.
- Leveraging and harnessing information streams. Content can serve as a catalyst to generate data, and that data can be further contextualized to create a new information stream. The content of a health-monitoring app is valuable in that it guides a user in how to use the app, but then the user-entered data can add further value, as the app analyzes that data and serves it back up to the user (Ismail, 2020).

Automate Continuous Delivery Pipelines

Authors working within Agile environments (mentioned in chapter 2 of this collection) are likely aware of how a continuous delivery pipeline works. Rather than creating content during product development and delivering a comprehensive set of documents to be published with the product release, content is created and released sprint by sprint, alongside the product code.

The benefit of content availability is not limited to allowing the automation of a delivery pipeline. Content consumers have come to expect content that is current, accurate, in the right language, refers to the right geographic area, and meets any accessibility needs they may have. Being able to personalize content by all of these criteria, simultaneously and automatically, is no longer a luxury. Doing this for product content can be particularly complex, and a robust and well-planned operating model is the only way to enable personalization at scale.

Improve Collaboration across Value Streams

First, content must be recognized as a value stream, and then it must have the ability to be used across an enterprise without obstacles. This leads to content interoperability, which can be defined as the ability of a range of systems, devices, and applications to access, exchange, integrate, and use content—a definition adapted from the Healthcare Information and Management Systems Society (HIMSS, 2021).

This principle dictates that interoperability must be possible in an automated way, without manual intervention or remediation, within and across content assets, departments, products, and organizations, with seamless portability.

Make Semantics the Foundation

Information science recognizes three main types of metadata: administrative, structural, and descriptive. We list them here with brief definitions taken from *Metadata Matters* by John Horodyski.

- Administrative metadata “provides information that helps manage an asset. Two common subsets of administrative data are rights management metadata, which deals with intellectual property rights, and preservation metadata, which contains information needed to archive and preserve a resource.”
- Structural metadata “indicates how compound objects are put together—for example, how a digital image is configured as provided in Exchangeable Image File (EXIF) data, or how pages are ordered to form chapters.”
- Descriptive metadata “describes an asset, an object, an item for discovery and identification as you would do in a search on Google or any other search tool. It includes elements such as title, creator, author, and keywords.”
(Horodyski, 2022, p. 16)

This classification of metadata, however, doesn’t tell the whole story. The three types of metadata work together to create possibilities beyond what any one of them can allow by itself. The three forms of information enablement absolutely depend on content being semantic. Semantically rich content “is content to which we’ve added machine-readable information that describes the content. Semantic information can be conveyed through both structure and metadata” (Rockley et al., 2015, p. 110). What allows the programmatic processing to show the right content to a particular user at a particular time, and to do so at scale, is the existence of applications that read and process any metadata attached to a piece of content. The more metadata is attached to a piece of content, the more likely that piece of content can surface at the appropriate time.

- Structural metadata reveals the relationship between elements within a single piece of content—and relationships between pieces of content. For example, the first paragraph of an article or page may have metadata indicating that it is the heading or title; the second paragraph may have metadata indicating that it is a description. Structural metadata might indicate a list or list items, a footnote, and so on. This is a good first step toward finding the right piece of content from a large repository, but more semantics are needed to help pull out a specific piece of content.

- Categorization semantics help determine intent. In other words, why would a user want to read this piece of content? This can determine the difference, for example, between looking for a pizzeria to dine at, buying a pizza pan or stone, finding a recipe for pizza, or learning about where pizza was invented.

Metadata can be added at different stages of the production process, from before a piece of content is created, during the authoring process, down the line in the supply chain, to when the content is being delivered to the presentation layer. The use of metadata is what allows us to plan for change. As new contexts reveal themselves, as new devices are brought into the mix, as new applications are developed, the content can be ready. Metadata helps future-proof ContentOps.

Consider Technical Capabilities

We can describe technical capability as having sufficient tools, technologies, and skills to allow an enterprise to operate with efficiency. Without a robust, fit-for-purpose set of tools, the ability to operationalize content production can be seriously hampered. Sadly, this is the case all too often. Content producers, from subject matter expert authors to technical communicators, are provided with software—such as word processors or software development platforms or databases disguised as spreadsheets—and expected to create kinetic content.

In the context of ContentOps, technical capability means using production-grade software, with customized interfaces, that allows content producers to move efficiently down the content supply chain from authoring to delivery. This also means being able to manipulate content with the granularity needed, which could mean supporting significant amounts of variants, reuse, and adaptiveness.

Some of the questions to pose when determining whether a content team has the right technology in place are whether they can accomplish the following:

- Automate portions of authoring, through transclusions, AI, or other means.
- Add metadata to content categories during the planning process.
- Create modular content components.
- Have the software validate their content structure as they write.
- Categorize content with metadata during the authoring process.
- Support authors to optimize content quality.
- Transfer content to and from a translation management system.
- Track localized content back to the source language.
- Track and record events during workflow for audit purposes.
- Assign metadata to images, including alt text for accessibility purposes.

- Aggregate content components into larger content documents.
- Eliminate manual steps by automating these tasks with software.
- Rely on automated delivery of content to its end point.

Conventions vs. Standards

When creating content for a level of robustness where it could be combined, recombined, and aggregated for multiple outputs, there is a need for tight conventions and standards. Search engine results reveal that these terms tend to be used interchangeably. However, as Rahel has taught for many years now, there are critical differences between the two concepts, and confusing them can be the difference between contributing to a successful model or having one that is subpar.

Editorial Conventions

Conventions can be thought of as editorial patterns that contribute to human comprehension. However, conventions are not hard-and-fast rules. Conventions can be shifted as needed to conform to a particular genre. Table 4.1 shows an example in which conventions help us understand the structural elements of a document, which are readily identified by virtue of their placement in relation to other elements.

Table 4.1: Example of an information asset with typical conventions applied.

Structural elements in the example	
Header	
Footer	
Page number	
Main heading	
Subheadings	
Text	
General (unordered) list	
Instructions (ordered) list	
Tables	
Table headings	
Images	
Image captions	
Footnotes	

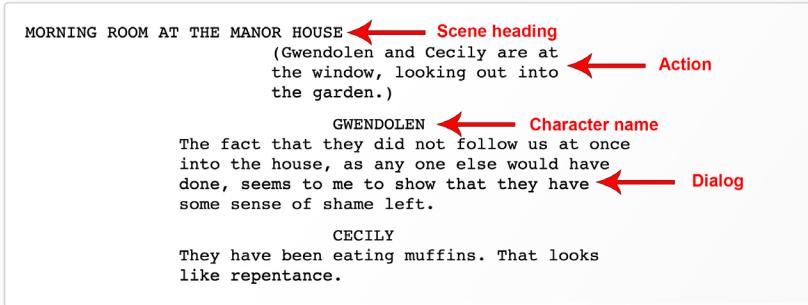


Figure 4.1: Sample of conventions in the screenwriting format. Credit: Carlos Evia, with content from Oscar Wilde's *The Importance of Being Earnest*.

This specific convention applies to Latin script languages and other left-to-right languages. Left-to-right languages, such as Hebrew or Arabic, will have different conventions, as will logographic languages such as Chinese, alphabetic syllabary languages such as Korean, and languages with a range of other types of scripts.

Figure 4.1 shows an example of a convention specific to scriptwriting. There are prescribed formatting conventions for elements such as scene headings, actions, character names, and dialogue. The rule about left-aligned text helping comprehension does not apply in this industry. Scriptwriters and readers are accustomed to the formatting and know how to interpret it.

Some conventions are not based on a notation system or genre but have become codified arbitrarily along geographic or cultural lines. For example, book titles on spines follow local conventions. Books in the UK, US, English Canada, Scandinavia, the Netherlands, and the Commonwealth have spine titles that are right side up when books are stacked face up. Most of continental Europe, Latin America, and French Canada have spine titles that are right side up when the books are stacked face down. In languages with Chinese-influenced writing systems, the title is written top-to-bottom.

Breaking with convention does not impede comprehension. For example, the advertising copy for a hamburger in Figure 4.2 does not follow a standard or even a convention. The usual order of headline, slogan, sub-headline, and body copy has been reordered for stronger reader impact. The order of the elements may be unique to this advertisement, yet people can understand the message. This would not be the case on the technical side, as we will discuss in the next section.

Technical Standards

A standard, according to the British Standards Institute, is an agreed-on way of doing something. It could be about making a product, managing a process,



Figure 4.2: Sample of an advertisement breaking conventions. Credit: StudiousGuy: Advertising Copy: Definition, Types, Examples.

delivering a service, or supplying materials. Standards can cover a huge range of activities undertaken by organizations and used by their customers. Standards are the distilled wisdom of people with expertise in their subject matter and who know the needs of the organizations they represent—people such as manufacturers, sellers, buyers, customers, trade associations, users, or regulators (BSI, n.d.).

The best-known standards body is the ISO (International Organization for Standardization), with close to 170 national standards bodies that have developed over 21,000 international standards. Groups of experts in a particular topic work together to ensure that everything from units of measurement to quality to video coding works seamlessly together by conforming to a particular standard.

There are many other standards bodies. For example, ANSI (the American National Standards Institute) administers and coordinates the US voluntary standards and conformity assessment system. AIM—a global industry alliance for stakeholders—covers radio-frequency identification (RFID), barcoding,

smart devices, and other automatic identification and data capture (AIDC) technologies supporting blockchain, Internet of Things, and real-time locating applications. The IEEE SA (the Institute of Electrical and Electronics Engineers Standards Association) develops standards across a range of industries. This is what ensures that electrical and digital devices can interoperate seamlessly.

The following are the standards bodies with the most relevance to content production.

- OASIS Open—Organization for the Advancement of Structured Information Standards. Produces open standards for SGML, XML, DITA, DocBook, XLIFF, and CMIS, among another thirty standards. The standards mentioned here fall squarely within the realm of Content Operations. Many of these are markup languages that make it easier for content producers to manage content—more efficient authoring, editing, tagging with metadata, reusing, revising, and generating multiple outputs. Other standards help with the seamless transfer of content between systems, whether between an authoring system and translation management system, between content management systems, or between other types of systems.
- W3C—the World Wide Web Consortium Open Web Platform. Produces standards for application development. This includes standards for building and rendering web pages, including HTML, CSS, SVG, Ajax, and other technologies for web applications. HTML is the best-known markup standard for content.
- Ecma International—originally the European Computer Manufacturers Association. Founded in 1961 to standardize computer systems in Europe, it now involves companies worldwide that produce, market, or develop computer or communication systems. JSON, Office Open XML, and Java are Ecma standards.

Standards, whether for devices or content, ensure interoperability. The reason you can plug an appliance into an outlet and be confident that the appliance will work is because of a range of standards that work together. In fact, the reason that you can be confident that the connector on an appliance can be plugged into an outlet is because of standards.

In the content realm, standards ensure that content is delivered faithfully, seamlessly, and automatically. Standards do not affect editorial conventions, but they do affect how content producers create content.

How Conventions and Standards Work Together

Combining technical standards with editorial conventions is the right move to ensure that systems will deliver content to its endpoint accurately, and, once the content gets there, that it makes sense to the readers. For example, compare

Table 4.2: Comparison of elements from two content types.

Editorial elements of a news release	Editorial elements of an event
Headline	Title
Status (for general release or under embargo)	Address
Text	Text
City (dateline)	City
Date and time	Date and time
Summary	Summary
Subtitle	Subtitle
Text	Text
Image	Image
Contact	Contact

the editorial elements of a news release and an event (Table 4.2), both of which have the same elements, aside from one genre-specific element (status vs. address).

When content management systems are faced with these two technical content types, or schemas, the systems can parse the elements and process the content quite easily. Applying a different formatting style sheet to each of the published outputs could format the content to look very different—the news release could be very plain, while the announcement of an event could look more like an invitation. However, that does not tell the whole story of what a computer can understand and what humans can understand. What makes the difference to a person is the editorial content within each of these elements. Even if one were to assume that the two content assets were from the same company and shared a single voice, the tone would be different. The news release will have a particular tone appropriate to the genre; the event announcement will also have a genre-appropriate tone.

In the example from Table 4.2, the way that content is constructed is relatively straightforward. If the content were to be single sourced, reusing all of the common elements, then creating it could be a little more challenging. In the following section, we will focus on the challenges involved in creating content for a complex environment—an operating model that needs content to be automatically discoverable, reusable, reconfigurable, and adaptable.

Putting Editorial Conventions to Use

Editorial conventions fall into a couple of subcategories: structure-related and copywriting-related. Both work together within an operating model that allows

for efficiency in content production and delivery—and effectiveness in communication and comprehension.

When analyzing content from an operational point of view, it is helpful to distinguish between levels of complexity when dealing with the written word. Considering the differences between concepts may seem pedantic now, but it can make the setup and configuration of content much easier when, for example, a supervisor or client requires more personalization, or modeling for omnichannel distribution, or granularized delivery to downstream systems.

- Copy, as carried out by a copywriter, is the act of producing words. Some definitions limit copywriting to the genres of advertising or promotional copy, but in the wider context of content professionals, copy refers to the wordsmithing aspect of writing. There may be some attention paid to editorial conventions but not to the technical aspects that improve ContentOps.
- Metadata is “data about data,” but its self-referential nature does not illuminate its importance in the realm of content production. Adding metadata to copy provides critical functionality that allows computer systems to understand how to process content. Metadata is part of the larger field of semantics—using signs and symbols to derive intent, most commonly for search but also for many other contexts.
- Content is the substance within a container, such as the material between a pair of tags. It can also be defined as contextualized data. As Carlos alluded to in the introduction to this volume, in the Content Operations arena, these definitions can be combined to form the following equation:

Copy + Metadata = Content

To bring this concept to life, consider the following example, popularized by Rahel. An extremely simple piece of copy is written: 12. As there is no context, neither person nor machine would understand its significance. As soon as the metadata of “month” is added (<month>12</month>), a computer system could interpret 12 to mean the twelfth month and display the results accordingly. A person reading “December,” or even “12th month,” can then derive meaning that would not be possible without the metadata.

Content asset is a collection of content recognizable as a particular type. Common examples are outputs, such as a printed document or a web page, or by genre, such as recipe or an event. Before the web, the most common example was a document, whether a single page or a full book. Today, a content asset can range from a single word or symbol, which could be a label on the button of a user interface, to a collection of content that becomes voice command routines for controlled virtual assistants.

Taking this distinction a step further is the addition of metadata to structure as well as copy. This type of content is called *intelligent content*, and it is the key to a high-functioning operating model.

Principles of Intelligent Content

The principles of intelligent content, as articulated by Rockley et al. (2015) in *Intelligent Content: A Primer*, have become accepted by content professionals across diverse industries. These principles are the foundation of content ecosystems where considerations such as personalization and publishing at scale are paramount.

Intelligent content is structurally rich and semantically categorized. The benefits of intelligent content are that the content becomes discoverable, reusable, reconfigurable, and adaptable. In practical terms, consider what this means:

- Structurally rich. Content is structured in a way that computers can understand how to use it.
- Semantically categorized. The content has metadata that gives the content container more meaning, so that the content can be processed with more specificity.

These dual aspects of intelligent content assist a content team in four ways.

1. Content becomes automatically discoverable. Search engines, whether within an application, on-site, or publicly available ones, can find the right content because they understand not only the words but their intended meaning. The function could mean something as straightforward as a CMS rule to “find any content with a particular tag” (for example, “troubleshooting”) and display it on a page about troubleshooting. An example from the world of e-commerce is the method for distinguishing multiple products with the same name. Multiple products called “flags” could have metadata that disambiguates an advertising flag from a national flag from a Post-It tape flag.
2. Content becomes reusable. A particular piece of content can be written once and reused many times, whether across multiple information products or within the same piece of information. The content, with its semantic categorization, is stored, and every time that piece of content is needed, the system issues an instruction for that content, tagged with that particular metadata, to be included in the right place within a content asset. An example might be reusing the phrase, “Be sure to use non-scratch utensils only.” Being able to write it once and use it thousands of times across recipes on a website saves time in writing (and later, in possible editing); promotes consistency; and drives significant savings at the time of translation.

3. Content becomes reconfigurable. The content can be aggregated in multiple ways, without manual intervention, to meet a new need. Adding enough semantics at the time of content creation means that the content has enough future-proofing to allow its use in new contexts with minimum fuss. A good example is that structurally rich information, such as flight, hotel, and car rental reservations, can be used in aggregation sites to do product comparisons but can also be used to set up a travel management instance on a service such as TripIt, or be pulled into your calendar for easy viewing, or be combined into some new information product that is yet to be invented.
4. Content becomes adaptable. The content can change its look and feel automatically, because the formatting is kept separate from the content itself. In other words, when delivering content to a specific channel, the content can be formatted automatically according to the styling rules for the size and profile of the device. Compare how LinkedIn articles, for example, are displayed on a mobile phone versus on a laptop computer.

Planning for an Operating Model Based on Intelligent Content

The first phase in a content life cycle is planning, and planning has never been more important than when preparing content for a production pipeline in a high-efficiency operating model. The structural metadata and semantic categorization need to cover at least three basic life cycles—the customer life cycle, the product life cycle, and the content life cycle—and plan for multiple output variants. The following example demonstrates the complexity of setting up a content ecosystem to enable the right content to be written and tagged for automated delivery by the systems at the publication end.

By mapping each touchpoint on the customer life cycle to the product life cycle, ContentOps is improved because content creators have identified the following:

- Which content needs to be created for each touchpoint, ensuring that there are no content gaps as the customer moves along their journey.
- Which content is needed for each stage of the product life cycle, from the time a product is launched until it is retired and replaced with a new version.

Each piece of content can be turned into intelligent content as it is created, which allows it to be surfaced to customers at the appropriate moment in their journeys. For example, a customer looking to get a software license for a new digital product will get a different message than a customer who wants to renew a license to a similar product that is about to be retired.

Each piece of content can also be categorized by a number of other factors: from the device and platform being used—think of slight variants between instructions for an iPhone vs. an Android phone—or by level of expertise, by geographic market, by language, and so on. Adding metadata provides the many benefits of intelligent content.

If a company were to create unique content assets for each variant shown in Table 4.4, that would result in an exponential proliferation of content that would then need to be maintained. Changes are made to a product (or to how a feature works) relatively frequently, and that has a tremendous impact on maintenance of multiple pieces of content. Using intelligent content principles, the proliferation of source content can be kept to a minimum. This means that instead of creating a separate piece of content for each variable, as shown in Table 4.4, a single piece of source content could include specific content components, shown in Table 4.3. For example, if content is needed in the Discover

Table 4.3: Potential content assets, by stage in the user journey, needed for a single product or service.

STAGES	Discover	Inform	Consider	Purchase	Support	Renew
Research	X					
Introduce		X	X	X		
Develop		X	X	X	X	
Evaluate				X	X	X
Version					X	X
Sunset					X	

Table 4.4: Potential content variants needed for a single product or service.

Device	Audience	Market	Language	Version	Product variant	Platform
Laptop	Novice	UK	UK English	1.0	Free version	iOS
Tablet	Intermediate	USA	US English	1.1	Basic access	Android
Mobile	Expert	Canada	AU English	2.0	Classic pricing	Windows
Wearable	Administrative	Australia	CA English	3.0	Pre-use pricing	Linux
Kiosk	_____	New Zealand	CA French		Premium access	
Bot	New grad	India	CH French		Enterprise access	
	Young family					
	Middle-aged					
	Seniors					

stage, that single piece of content could include variants to cover all device types, audiences, markets, and so on.

A content team can use editorial techniques to support ContentOps in at least two ways. One way involves structure; the other involves writing. These two aspects work together to create fit-for-purpose content that works to deliver into multiple outputs and to multiple audiences.

Applying Editorial Structure to Copy

The structural considerations for creating content that works within a complex content ecosystem, where automation for scale is important, can be linked to the design philosophy of minimalism. The concepts associated with minimalism were developed in the 1980s and 1990s by John Carroll and his colleagues at IBM, and they led to modularity—the chunking of content into modules. Carroll’s publications, including “The Minimal Manual” (1987), *The Nurnberg Funnel* (1990), and *Minimalism beyond the Nurnberg Funnel* (2003), demonstrate how minimalism is anchored in instructional design principles, making it easier for adults to absorb practical information.

Minimalism

The minimalist approach to designing instruction and documentation combined cognitive psychology and instructional design to look at the science behind how people absorb information. The approach relies on task orientation to produce more effective learning and more rapid performance outcomes.

The four main design principles of minimalism, as summarized by Hans van der Meij and Carroll (1995), are:

- Choose an action-oriented approach. Users typically want to do things. This principle reflects the user-centeredness of minimalism.
- Anchor the tool in the task domain. A tool is a means to an end. This principle asks designers to select training tasks that are meaningful for the user.
- Support error recognition and recovery. To err is human. There are several ways to increase user competence and comfort levels in handling mistakes.
- Support reading to do, study, and locate. Designs must fit as much as possible the diverging needs and propensities of the intended audience. This principle reflects the user-centeredness of minimalism.

Minimalism dictates that authors must know their users and their needs and can understand their goals, tasks, and interactions. Content is used out of “sequence,” and in that regard, there is a need for content that can be

automatically pulled, according to the requirements of the moment, and delivered into multiple contexts. This becomes the backbone of ContentOps.

Modularity

If minimalism is a design philosophy, then modular content is the embodiment of that philosophy. Well-written modular content works somewhat like Lego blocks. The blocks are not the same size or color, but they are manufactured in predictable shapes and sizes, which makes it possible to build larger structures. A simple example is shown in Figure 4.3.

For modular content to work within a scalable ecosystem, it must have the following characteristics:

- A module must be a discrete, self-contained, reusable unit of information. One recommendation is that the size of the unit should be the smallest logical standalone piece of information. In some cases, this could be a sentence, a paragraph, or a larger collection of elements that work together as a cohesive unit.
- A module must have enough associated metadata to allow it to be pulled by a system in a way that ensures that content is shown at the right time, in the right place, to the right audience. This is the mechanism behind personalization, among other purposes.
- Each module must follow a common schema. This allows systems to process content with reliability, as the structural metadata indicates content type, purpose, and intent to those systems.

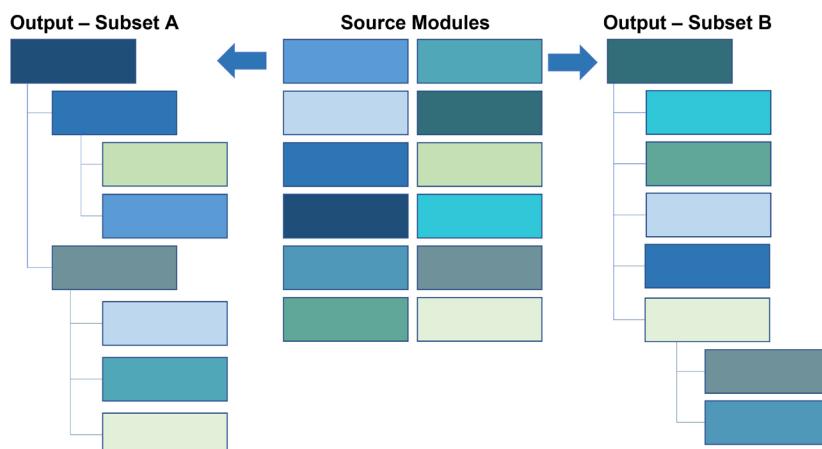


Figure 4.3: Content outputs from a single source of content modules. Credit: Rahel Bailie.

- Each module must be portable, so that the modules can be mixed and matched with other modules to build a creative whole. Each module (and all of them as a whole) must maintain a level of coherence for the content's purpose to avoid remixing content just for the sake of remixing. Metadata can help with this process of assemblage and reconfiguration based on audience and purpose.

Understanding how modular content can be is a key part of understanding how to create an efficient system for ContentOps.

Applying Writing Techniques

Learning about writing techniques is an ever-changing, never-ending journey. As norms change, both in reader expectations and the demands of products and technologies, content teams will be challenged to determine which techniques are most appropriate to their ContentOps. Having said that, operating models can be quite varied; there will be a need to continually evaluate the appropriateness of writing techniques to feed the content ecosystem.

Five core editorial practices are fundamental to working within a high-efficiency operating model. The paradox of the advice to adopt these five practices is the caveat that comes with that advice: don't fully adopt these five practices. Sometimes, the problem will be that a principle of one practice contradicts the principles of another practice. At other times, the problem will be that the principles to make copy more friendly to read can backfire when used in a complex content ecosystem where contexts can change, depending on the reuse scenario. Combine the best and most appropriate practices to suit a particular operating model.

1. Plain Language

The Plain English movement began in the UK in the mid-twentieth century as a way to help the average resident complete overly complicated government forms. Today, most English-speaking countries have a government mandate to produce content in Plain English. Given the number of languages and the range of grammatical constructions, there is no single right or wrong way to go about writing in plain language—though for the audience of this book, the source language will overwhelmingly be English. Instead, adopt and adapt: adopt the ethos of the practice and adapt it to the languages used by your organization.

The principles of plain language are designed to allow readers to concentrate on the message rather than be distracted by complicated language. These principles were inspired by and adapted from the work of Robert Eagleson (1990). The ethos

of plain language is to “keep it simple without dumbing down the message,” and the guidelines for writing in plain language copy can be summarized as follows.

- Write using active voice to make the meaning clearer, and use first-person, informal address when possible.
- Break complex sentences into smaller chunks. Limit each sentence to a single concept. Don’t break a concept into multiple sentences.
- Compose short sentences with simple structures. Use a subject-verb-object structure. Keep readability appropriate for “skip, skim, scan” reading patterns.
- Use common words—and fewer words when possible.
- Edit out irrelevant information.
- Avoid idioms and fuzzy language.
- Agree with a structure to set context involving the author and the reader.
- Keep the message brief and focused.

There are several benefits of using plain language writing techniques. Plain language is easier to read for all reading levels. Readers with low literacy skills will have an easier time reading, and readers with higher literacy skills will find it easier to skim the content. It’s also harder for inconsistencies and inaccuracies to creep into plain language. Passive voice can hide who is taking an action, whereas active voice sentence structure promotes clarity.

In a complex content ecosystem, certain aspects of plain language are not practical to implement. For example, using pronouns such as “it” or “their” can have unintended consequences when content is reused. When the pronoun begins a piece of content that follows a new piece of content, the pronoun may make an inappropriate reference. Given that a high percentage of readers don’t understand the rules for pronoun reference, using the noun instead of the pronoun is recommended.

2. Controlled Language

Controlled natural language works hand-in-hand with plain language and global readiness. Controlled vocabularies are often specific to an industry, such as aviation or medicine.

Simplified English is a predetermined set of words and writing rules specifically made to help write instructions and instruction manuals so that any individual in the world can clearly understand the manual. It is especially helpful for those whose native language isn’t English. Simplified Technical English (ASD-STE100, 2021) is likewise a predetermined collection of words and writing rules, classified as a “controlled language,” that was initially developed by the aerospace industry and later adopted by other industries. Controlled natural language focuses on simplifying technical instructions.

Common writing rules for controlled natural language include the following.

- Having a dictionary of acceptable terms, to force consistency and decrease comprehension errors.
- Setting only one meaning and word form for each word. This means that a word such as “list” would have one meaning, either tilting of a boat or enumerating items, and if used as a verb, it could not also be used as a noun. To reduce complexity, choosing the simplest term possible is encouraged. Homonyms and polysemous (words with multiple meanings) words are forbidden. Because controlled language is about precision, the use of pronouns is eliminated, as is the use of synonyms.

In an operating model where controlled natural language would help with content production efficiencies at scale, then likely content optimization software would be used to automate terminology and grammatical compliance.

3. Global Readiness

Your content may not have a global audience—or, should we say, not that you know of. Even if your audience is a single country or market, chances are that the audience is not homogeneous. Some of the audience may have English as a second (or third, or fourth) language or be used to a different type of English. Others may understand English well enough but struggle with idioms or cultural memes.¹

Global readiness affects not only content—what is displayed and its metadata—but also the taxonomy, which can feed a navigation scheme; image captions and callouts; the titling, captions, and metadata behind it; podcast titling and transcription; and transactional copy, such as navigation and interface text, labels, and help, error, and administrative messages. Val Swisher, in books such as *Global Content Strategy: A Primer* (2014) and *The Personalization Paradox* (2021), provides guidelines for writing for a global audience, such as using present tense in sentence construction; avoiding passive voice; and avoiding abbreviations, noun strings, and gerunds. Likewise, Loy Searle analyzes the concept of localization (and related practices) for ContentOps in chapter 7 of the present volume.

The considerations are many, and in some ways, they are a moving target—the fluidity of language could mean that a culturally sensitive image this year could be a problem next year. It is important to understand how global readiness for a body of content can affect ContentOps. Translating content with minimal human intervention, producing alternate-language content that follows corporate guidelines, and meeting regulatory requirements are all operational considerations. (The idea that consumers expect brands to have good, usable content may be outside of the scope of Content Operations but is certainly a nontrivial issue.)

¹ Vivid examples can be found on YouTube by searching “Complete the Meme,” “Finish the Meme,” and “Strange English Idioms.”

4. Accessibility

Keeping content accessible warrants discussion, because many countries require organizations to meet accessibility guidelines. For example, the UK and US governments have comprehensive guidelines for accessibility. Because class-action lawsuits based on lack of accessibility are steadily increasing, attending to accessibility is also seen as a risk-management strategy. Producing accessible content should not be seen simply as a business obligation—though some companies are more apt to pursue accessibility when they learn that accessibility techniques are also beneficial for search engine optimization—but also as a way to respect end users, in the way that user experience is seen as a way to serve those end users' needs.

Some accessibility requirements fall within the purview of the technologists, such as ensuring that a blink rate does not trigger epileptic episodes. Other accessibility requirements are shared between technologists and content professionals. For example, technologists ensure that alt-text fields are made available, and content professionals ensure that they write text that is both accessible and meaningful for readers.

How content is set up for reading makes a large contribution to accessibility. The following guidelines are meant to be used in conjunction with the plain language, controlled vocabulary, and global readiness already discussed.

- Multiple access methods
 - Provide multiple forms of content; do not rely on voice-only or text-only methods of accessing content.
 - Consider live interpretation as an option for live content delivery.
- Unique page titles
 - Identify the subject of the topic or page with the title.
 - Front-load titles to include important words at the beginning.
 - Use hyphens as separators for easier reading; do not use ampersands.
- Section headings
 - Break text up into logical chunks with headings.
 - Use style sheets to add headings (not in-line formatting).
 - Give the headings meaningful descriptions.
- Lists and tables
 - Format lists and nested lists using structure tags.
 - Use tables only for tabular data.
 - Ensure that tables can be read properly by screen readers.
- Links and alt text
 - Use unique descriptive links for hyperlinks and cross-references.
 - Use canonical links for websites (and test them with a screen reader).
 - Ensure that alt text describes visuals in a way that is meaningful to the reader.

- Avoid in-line links that interrupt reading; place links at the end of a paragraph or section.
- Write descriptive links to downloadable files.

The important thing to remember is that content needs to be accessible across the board, not only for users with one type of impairment. There are readily available accessibility resources for just about any content question you may have, whether thinking about websites, video games, content hubs, technical documentation, or live presentations.

5. Diversity, Equity, and Inclusion

In the book *From Solo to Scaled*, Natalie Marie Dunbar (2022) includes a series of recommendations for building a successful content practice. Dunbar urges her readers to “embrace the intention of inclusive design: to include everyone, regardless of differences in ability, gender, race, sexual orientation, or other characteristics, in the activities of the practice” (p. 134). Dunbar cites the following definitions from Project Inkblot (2019), which, coincidentally, Carlos has been using in his introductory college courses on topics related to Content-Ops after a recommendation from Sarah O’Keefe.

Diversity is *quantitative*. It’s the composition of different people represented in what you make and the decision makers on your team.

Inclusion speaks to the *quality of the experience you’ve designed* for these diverse folks, so they experience themselves as leaders and decision makers.

Equity lives in how we design our systems and processes; the way we work, and who we work with, so we are upholding our commitment to diversity and inclusion.

Dunbar adds that building diversity into a content practice “asks you to think about who is *not* represented in your products and services, and what it will take to include them—a process that often begins with how you ‘speak’ to users with content and how your tone shifts to be inclusive for all” (p. 135).

Although there are no global standards for incorporating diversity, equity, and inclusion (DEI) into writing and content—and those vary considerably across locales, contexts, and languages—some authors provide specific recommendations. For example, Michael Metts and Andy Welfle (2020), in *Writing Is Designing*, introduce the following guidance on how to write more inclusively:

- Avoid language that assigns value to traits (e.g., say “she uses a wheelchair” instead of “she’s confined to a wheelchair”).

- Avoid prescriptive language when talking about people (e.g., say “this software helps you create graphics, photographs, and illustrations” instead of “this software is for graphic designers, photographers, and illustrators”).
- Adopt the singular “they” (e.g., say “they will be in touch with you” instead of “[s]he will be in touch with you”).
- Read more about inclusivity.

Talking about DEI and ContentOps is important, but it is also easy to fall into the traps of stereotyping or infantilizing traditionally misrepresented communities. Therefore, we recommend involving DEI experts as consultants or, if possible, bringing them on as full-time members of a content team. (In that spirit, we invited Jonathan McFadden to write the epilogue of this collection as an interplay between content and DEI.)

Content Delivery

In chapter 8 of this collection, “The Technology that Supports Content Operations,” Patrick Bosek discusses tools and processes related to the processing and delivery of content in complex operational environments. In this chapter, and particularly in the next section, we want to emphasize that there is a reason for investing in robust ContentOps models and workflows. Actually, there are several reasons that relate to saving time for content teams, improving content accuracy (by reducing dependency on manual processes of copying and pasting, for example), allowing growth and scale of projects, reducing risk in regulated industries, improving time to market of information products, and increasing trust along the customer journey. Patrick’s chapter provides details about software and apps that process and deliver content. Here, we want to focus on what can be described as the North Star of robust ContentOps: offering content as a service instead of publishing it.

Content as a Service

Content as a Service (CaaS) is about establishing interaction frameworks between people, organizations, devices, and services so that there is accountability for the resulting behavior and, presumably, a measure of effectiveness when viewed from the perspective of all legitimate stakeholders. This adapts some of the work by Joe Gollner (2016), who has made a habit of working with complex content projects. When exploring this level of content flexibility, Gollner likes to talk about “authority networks” that link a series of actors leading up to events such that there is traceability from observable events back to the responsible parties.

In order to reach the North Star of CaaS, content should have the following characteristics.

- Profiled. The content has been semantically structured in a way that suits delivery for personalized use. For example, a piece of content could be tagged with metadata so that it would be delivered to customers who bought a particular version of a product, arriving at the time that their product warranty is about to expire.
- Offered. Rather than publishing content, content is delivered to a neutral location, such as a database, and later supplied for use by multiple technologies. For example, a digital health monitor does not know what the health level of its user will be on a particular day, so all potential answers will be stored, with the appropriate answer delivered depending on the situation on any given day.
- Dynamic. The content, at its source, can be updated as needed and continuously delivered to all potential end points. For example, product content is being continuously created and updated throughout the development phase, so that it is completed at the same time as the product.
- Independent. The content is stored as standalone, autonomous units that can be used in multiple contexts. Contexts could mean devices, products, audiences, and so on—or a combination of contexts.
- Ubiquitous. The content is online, searchable, and findable.
- Molecular. The content is not formatted as specific documents. Instead, it is presented as information molecules.
- Spontaneous. The content will be triggered by contexts.

Gollner adds that content is planned, designed, created, managed, and exchanged as objects that incorporate not only the editorial side of the content but also the associated rules governing the structure and meaning of the content. This means that the content has an array of rendition and behavior processes that a module can use to render that material independently or in concert with other content objects. In practical terms, we can see that the content being managed is highly precise, so that a variety of application processes, including multiple rendition scenarios, can operate on that content with a high degree of confidence and effectiveness. Content coexists with the complete array of known behaviors that it supports, and it is exchanged with stakeholders who will deliver and use the content and behavior in their own environments to meet their own goals.

Now that we have established the foundations in principles of kinetic content, the interplay of editorial conventions and technical standards, the structure-based benefits of intelligent content, and the basics of dynamic processing and delivery, in the next section we introduce a model for determining the level of robustness of a ContentOps implementation.

Levels of Robustness for ContentOps

As we introduce our levels of robustness for ContentOps, we acknowledge that they are partially inspired by the levels for developing technical content that Alan Pringle and Sarah O'Keefe included in their book *Technical Writing 101*, which Carlos cites frequently and to this day includes in his university-level courses about content development and strategy. Alan and Sarah proposed the following levels for classifying a methodology for developing technical content:

- Chaos (“there’s no consistency in the presentation of content”)
- Page consistency (“content looks the same on paper [or other delivery format],” but there’s no consistency in the source files)
- Template-based authoring (content follows “predetermined styles [and] writers don’t spend time figuring out how to create particular formatting—they apply styles to add formatting”)
- Structured authoring (“a publishing workflow that defines and enforces consistent organization of content”) (Pringle & O’Keefe, 2009, pp. 41–42).

We also acknowledge the maturity model for Content Operations proposed by Colleen Jones (2019) to “help your company identify your current level of content operations and then decide whether that level will support your content vision and strategy” (p. 163). The model proposed by Jones includes the following levels:

1. Chaotic: “No formal content operations, only ad hoc approaches”
2. Piloting: “Trying content operations in certain areas, such as for a blog”
3. Scaling: “Expanding formal content operations across business functions”
4. Sustaining: “Solidifying and optimizing content operations across business functions”
5. Thriving: “Sustaining while also innovating and seeing return on investment” (Jones, 2019, p. 164).

Rachel McConnell (2022) mentions workflow, methodologies, and organizational alignment as some components of Content Operations. When describing the levels of robustness, we focus particularly on workflows and their affordances. Although not all workflows related to ContentOps involve writing, we base our working definition of the term on the description of a writing workflow as the process for completing a literate activity and the tools used in that process (Lockridge & Van Ittersum, 2020). Literate activity, as Prior defines it, involves “reading, talking, observing, acting, making, thinking, and feeling as well as transcribing words on paper” (1998, xi). We believe that this definition can expand beyond purely *literate* activities and cover processes and tools related to digital content creation for videos, images, maps, graphs, and the like.

We view the concept of affordance based on the following definition by Carolyn Miller: “A technological affordance, or a suite of affordances, is directional, it appeals to us by making some forms of communicative interaction possible or easy and others difficult or impossible, by leading us to engage in or attempt certain kinds of rhetorical action rather than others. Affordances both enable and constrain, they both pull on us and push at us” (Miller, 2010, p. x). If ContentOps were a video game, it would be more a sandbox game than a quest adventure. The levels of robustness would be equivalent to levels in the game: as the levels go up, so does their operational difficulty, but they also unlock new features and benefits (our affordances). Although Player 1 in a sandbox video game *could* advance to a higher level, that does not necessarily mean that they *have* to advance (as opposed to a quest, where there’s a hostage to rescue and/or a monster to defeat). The same applies to the levels of robustness in ContentOps: a company can be effective at a level that provides the right affordances for the company’s scope, audience, and interests. For example, if a small startup only generates content for a static website using a building system such as Square-space, there is no pressing need to advance to a higher level of robustness. Of course, if the scope changes and the content needs to scale (say, if the startup is acquired by a bigger tech firm), moving to a higher level will be a necessity, and the expansion can avoid painful points of friction by adopting the right workflows and tools even at a low level of ContentOps robustness. For this reason, we are intentionally not labeling the levels of robustness as a maturity model for Content Operations—which Jones (2019) has already proposed. Without getting into a deep criticism of “maturity model” as a construct, we agree with Verwijs (2019) in his assertion that “growth is not linear and it doesn’t happen in discrete phases marked by convenient external characteristics.”

The workflows included in the levels of robustness cannot avoid discussions of tools and technologies. As Patrick Bosek writes in chapter 8 of this volume, ContentOps could be separated from technology, but practically speaking, the two are deeply intertwined. Finally, we should mention that the levels are not absolute and that there’s always variance based on audience, purpose, and technological capabilities of any content project. Additionally, the workflow scenarios included in the levels are broad generalizations based on real cases (but fictionalized).

Level 0: Desktop Publishing at the Core of Any Content Work

Robustness is weak or nonexistent. Sarah O’Keefe has written that Content Operations “are not necessarily automated or efficient. If you move information from place to place via copy and paste, and have extensive manual quality checks in place to catch the inevitable errors, that’s still content ops (albeit inefficient and tedious content ops)” (O’Keefe, 2021). Content workflows at level 0 follow

some of those inefficient and tedious processes that probably can work for a small team with one single audience type or one channel for delivery—but most likely cannot scale to address the needs of additional audiences and channels.

Sample workflow: Authors create content in a word processor, which can be a standalone application on a computer (e.g., Microsoft Word) or a web app (e.g., Google Docs). A single content creator (or a small team) operating at this level may have a manual style guide in either a physically printed document or book or a desktop publishing file. Authors create content primarily for a single audience segment. If the information products created at this level need to address more than one audience type (e.g., novice and advanced users), authors have to create parallel sections in the document(s). The primary channel of delivery is a printed document or a PDF distributed online or through physical digital media. Authors depend on manual processes of copying and pasting for updating content, and there is no version control (other than the basic tracking features of the selected word processing tool at the center of this level). A new version of a document will replace (in some cases permanently) the previous version.

Affordances: Authors or small teams can control a few files directly in one single computer or a cloud-based drive. If using the desktop tool properly, authors can generate automated tables of contents that produce clickable outcomes in PDF exports. Any products developed at this level are static, and updates will require republishing.

Level 1: Content Moves Online

Robustness is weak but deliverables can be multichannel. Teams operating in level 0 can easily move to level 1 if their audiences' needs involve a basic website. This website can be created with a static HTML template or site builder or use a web content management system like WordPress, primarily as a copy-and-paste channel secondary to the core desktop publishing operation that ruled over level 0.

Sample workflow: Similar to level 0, authors in level 1 primarily create content in a word processing environment and may use a manual style guide for conventions. The primary channel for publication still is a printed document or a PDF export of this document. The copy-and-paste process in this level can include a secondary publishing channel for a basic website. This website can be based on a static HTML template or site builder platform, or it can be walled by a social media platform like Facebook or Twitter. Author(s) copy and paste sections of the word processing tool to an HTML template or to social media posts. In some cases, the team may have access to a web CMS (e.g., Drupal or WordPress) but use it mainly as a page-oriented copy-and-paste replica of their main desktop publishing document. Any updates to the original desktop

publishing source will not be reflected on the website/social media secondary channels unless authors perform a new copy and paste (and publish) update. Likewise, any updates performed directly on the secondary website/social media channels will not be reflected on the original desktop publishing source. Similar to level 0, any elements in the content that address different audiences will need parallel sections. Content teams (which can be very small or consist of one person) frequently mix content with interface considerations, and they might use a design tool like Figma for development that is locked in a single deliverable channel.

Affordances: Authors or small teams can address the needs of their users in a rudimentary multichannel approach with a printed/PDF deliverable and a website or social media deliverable. If using a web CMS, the team can gain basic insights from the tool's analytics. Teams may start looking at basic terms of content strategy and standards (primarily HTML) to manage the parallel publishing channels.

Level 2: Separation of Content from Presentation Opens Doors

Robustness is light. Level 2 is characterized by the adoption of tools and workflows that separate content from presentation. Whereas in levels 0 and 1 authors were primarily working with what-you-see-is-what-you-get desktop publishing tools and copy-and-paste exports of that content, in level 2, teams can start thinking about multichannel publishing and other advanced ContentOps features by embracing the separation of content from presentation. This separation, however, “can create philosophical and cognitive dissonance for technical communicators trained to think of information as content that is inherently linked to presentation” (Clark, 2007, p. 36). Mark Baker warned about the difficulties of teaching this concept: “One of the hardest things about moving technical writers from desktop publishing to structured writing is persuading them to give up responsibility for how the final output looks. Writers will keep looking for ways to specify layout, even in markup languages specifically designed to remove layout concerns. They understand their jobs in terms of the responsibilities their old tools imposed on them” (Baker, 2013, p. 87). Therefore, in level 2, authors and teams leave the page-oriented core of desktop publishing and focus on content that will be formatted at a later stage in a publishing process.

Sample workflow: Authors may work primarily on a web CMS or in a developer-friendly text editor. However, their authoring environment does not include decisions about layout or presentation. Authors create posts, pages, or—in some cases—modules of content that will be assembled and organized during an automated publishing process. In a CMS or CCMS (Component Content Management System), authors can work in interfaces that hide any code behind their writing. In a text editor, authors create content in

HTML code or, most likely, in a lightweight markup language like Markdown or reStructuredText (file formats that developer communities tend to treat as standards). Advanced level 2 teams can adopt conventions and standards (e.g., XML, JSON) and avoid copy-and-paste practices. Instead, the “raw” content can adopt different presentations as determined by the CMS. In some cases, the CMS allows exporting to PDF for print with automated formatting.

Affordances: The (primarily web) CMS or CCMS used in this level may have built-in version control features that allow authors to revert changes without losing information. Teams can exchange information without relying on copy-and-paste techniques. Scaling from a small web CMS to a more robust publishing tool is easier than coming from a desktop publishing application. Customization and filtering based on audience type or other criteria still require parallel sections. Web CMS tools at this level can have an automated feed in an XML grammar that enables basic syndication to other websites and even some social media platforms. At this level, content can be perceived as a valuable asset and not just a support feature for a product or service.

Level 3: Focus on Structured Content and Multichannel Delivery

Robustness is solid. From this level on, structured content is the foundation for robust ContentOps. The need for well-structured, semantically rich content that is human and machine readable is an enabler of the many forces that are shaping content conversations in industry and academia. Atherton and Hane define structured content as follows: “Structured content is content that is planned, developed, and connected outside an interface so that it’s ready for any interface. It allows you to define the skeleton of your content before you create it. Breaking content into the smallest pieces possible (within reason) so that it is free to go anywhere, anytime” (Atherton & Hane, 2018, p. 32). Level 3 is primarily defined by the adoption of structured authoring practices that move beyond the mainly presentational features of, say, tagging a section title with a heading 2 in HTML. Level 3 implementations are frequently driven by a tool or platform. An organization can involve content professionals in the selection of such tools, but the decisions mainly come from management or information technology departments. Tools in this level can include headless CMS platforms (content management systems that separate what the users see [the presentation layer] from what the developers and authors create and maintain [the backend layer]) with web and app deliverables as main channels of publication.

Sample workflow: The organization embraces content as a valuable asset, and the dedicated content team may expand to include authors and also strategist(s) and engineer(s). Authors mainly write content in the form of a page or post (with a website or app as publication targets), but they can write chunks of

content independent from a specific page that then will be aggregated with other chunks at the time of publishing. Behind the scenes, the CMS may structure content in JSON strings, HTML tags, XML structures, or a combination of those. Teams may follow a detailed content strategy with specific metrics for success, and they probably also adopt standards for style and development.

Affordances: In level 3, content is managed as computable data. This enables single sourcing of content that is stored in one repository and can be assembled on the fly, based on audience, purpose, and format criteria. Level 3 also enables personalization (more about this in chapters 5 and 6 of this volume) and content filtering for specific needs. In this level, content can be served as a service via an API (application programming interface) that does not require manual publishing from the authoring side.

Level 4: Highly Structured ContentOps

Robustness is appropriate for highly regulated industries. If level 3 is driven by a specific tool or content management platform, level 4 is driven by a content standard. Let us rephrase that: Level 4 is driven by the robustness needs (multiple audiences, multiple languages, multiple deliverables in a highly regulated industry and environment) that only a solid content standard for managing and publishing highly structured material can fulfill. One of the most popular standards in this category is the Darwin Information Typing Architecture (DITA). DITA consists of a set of design principles for creating “information-typed” modules at a topic level and for using that content in delivery modes such as online help and product support portals on the web (Day et al., 2001). Don Day—one of the original developers of the standard—explained that, when naming it, DITA “represented a great deal of messaging in a compact and memorable acronym”:

- Darwin: for specialization and how things could “evolve” from a base.
- Information Typing: for representation of knowledge as typed units.
- Architecture: a statement that this was not just a monolithic design but an extensible tool that could support many uses. (Day, quoted in DITAWriter, 2016)

We (Rahel and Carlos) actually met when we were members of the technical committee maintaining the DITA standard with the non-profit OASIS. However, when we decided to write this edited collection, we established the directive that we would not make it solely a mechanism for promoting the use of DITA over other alternatives that could be used at level 4. For example, many companies of diverse sizes have, with considerable limitations—particularly when it comes to managing multiple languages in a content

repository or word/phrase-level filtering—implemented and scaled Content Operations with JSON as their standard for managing and delivering structured content. Regardless of the standard, level 4 has the primary characteristic of treating content as computable data.

Sample workflow: Authors follow routines and experiences similar to those in level 3, but they also take advantage of a powerful standard to annotate and comment on the content with metadata and structural elements that will enable advanced filtering, automated translation and localization capabilities, and behind-the-scenes tags that can happen even at the level of a word or letter. Content is still separated from presentation, and the workflow allows for publishing or serving via an API with profiling and filtering. Examples include complex medical and pharmaceutical environments where a piece of content can be reused on a physician's tablet, a pharmacy's label, or the user interface of a laboratory machine with accuracy and real-time delivery, but can also filter through a plethora of variables (e.g., prescription dosage or collection of symptoms).

Affordances: More than in any other level, the affordances of level 4 focus on content reuse. In particular, “efficient content reuse does not involve copy-and-pasting; instead it uses transclusion, whereby content is authored in one location and used by reference in other locations” (Eberlein, 2016, p. 55). Content in level 4 is highly kinetic and adaptable, and its status as computable data enables deep personalization, translation, and localization, with capabilities for machine learning and artificial intelligence in processing and serving. Level 4 also includes the affordances unveiled in level 3, but with more capability for growth and development.

Conclusion

We mentioned earlier in this chapter that not all ContentOps require the same level of robustness. The mathematician Erica Jen explains that “presence or absence of robustness at one level does not imply presence or absence at another level, and perhaps the most interesting cases are those in which the interconnections among components not themselves robust give rise to robustness at the aggregate level” (Jen, 2003, p. 14). We see those interconnections as key to identifying gaps (and accomplishments) in ContentOps implementations.

In real life, however, content creation is more complicated than sticking to a recommended level of robustness or even adopting a specific technical standard or writing convention. In complex organizations, different teams could be operating at disparate levels of robustness (think of content marketing and technical communication with different tools and deliverables). Those areas of interconnection can allow organizations and teams to ventilate content silos and make smart decisions about required tools, training, and outcomes. However, regardless

of robustness, a ContentOps implementation would be nothing without solid audience consideration and research, and Kevin P. Nichols focuses on that in the following chapter, “Customer Experience and Content Operations.”

References

ASD-STE100. (2021, April 30). *ASD simplified technical English specification*. Retrieved November 18, 2022, from <https://asd-ste100.org/>.

Atherton, M., & Hane, C. (2018). *Designing connected content: Plan and model digital products for today and tomorrow*. San Francisco, CA: New Riders.

Baker, J.W., Schubert, M., & Faber, M.H. (2008). On the assessment of robustness. *Structural Safety* 30, 253–267.

Baker, M. (2013). *Every page is page one: Topic-based writing for technical communication and the web*. Laguna Hills, CA: XML Press.

BSI. (n.d.). *Information about standards—What is a standard?* Retrieved November 18, 2022, from <https://www.bsigroup.com/en-GB/standards/Information-about-standards/what-is-a-standard/>.

Carroll, J.M. (1990). *The Nurnberg funnel: Designing minimalist instruction for practical computer skill*. Cambridge, MA: M.I.T. Press.

Carroll, J.M. (Ed.). (2003). *Minimalism beyond the Nurnberg funnel*. Cambridge, MA: M.I.T. Press.

Carroll, J.M., et al. (1987). The minimal manual. *Human-Computer Interaction*, 3(2), 123–153.

Clark, D. (2016). Content strategy: An integrative literature review. *IEEE Transactions on Professional Communication*, 59(1), 7–23.

Day, D., et al. (2001, March 1). *Introduction to the Darwin Information Typing Architecture*. IBM developerWorks.

DITAWriter. (2016, July 15). *Don Day and Michael Priestly [sic] on the beginnings of DITA: Part 2*. Retrieved November 18, 2022, from <https://www.ditawriter.com/don-day-and-michael-priestly-on-the-beginnings-of-dita-part-2/>.

Dunbar, N.M. (2022). *From solo to scaled: Building a sustainable content strategy practice*. New York: Rosenfeld Media.

Eagleson, R.D., et al. (1990). *Writing in plain English*. Canberra: Australian Govt. Pub. Service.

Eberlein, K.J. (2016). Content reuse. In R. Gallon (Ed.), *The language of technical communication* (pp. 54–55). Laguna Hills, CA: XML Press.

Gollner, J. (2016, December 28). *Content 4.0*. Retrieved November 18, 2022, from https://www.gollner.ca/2016/12/content_4-0.html.

Gribble, S.D. (2001). Robustness in complex systems. In *Proceedings of the Eighth Workshop on Hot Topics in Operating Systems* (pp. 21–26). Elmau, Germany: IEEE.

HIMSS. (2021, August 25). *Interoperability in healthcare*. Healthcare Information and Management Systems Society. Retrieved November 18, 2022, from <https://www.himss.org/resources/interoperability-healthcare>.

Horodyski, J. (2022). *Metadata matters*. Boca Raton, FL: CRC Press.

Ismail, S. (2020). Information enablement and exponential growth. *Intercom*, 67(6), 6–10.

Jen, E. (2003). Stable or robust? What's the difference? *Complexity*, 8(3), 12–18.

Jones, C. (2019). *The content advantage [Clout 2.0]*. San Francisco, CA: New Riders.

Lockridge, T., & Van Ittersum, D. (2020). *Writing workflows: Beyond word processing*. Ann Arbor, MI: University of Michigan Press.

McConnell, R. (2022). *Leading content design*. New York: A Book Apart.

Metts, M.J., & Welfle, A. (2020). *Writing is designing: Words and the user experience*. New York: Rosenfeld Media.

Miller, C.R. (2010). Foreword. In S.A. Selber (Ed.), *Rhetorics and technologies: New directions in writing and communication* (pp. ix–xii). Columbia, SC: University of South Carolina Press.

O'Keefe, S. (2021, March 29). *Content operations (content ops)*. Scriptorium. Retrieved November 18, 2022, from <https://www.scriptorium.com/2021/03/content-operations-contentops/>.

Pringle, A., & O'Keefe, S. (2009). *Technical writing 101: A real-world guide to planning and writing technical content*. Durham, NC: Scriptorium Press.

Prior, P. (1998). *Writing/disciplinarity: A sociohistoric account of literate activity in the academy*. New York: Routledge.

Project Inkblot. (2019, September 18). *How to begin designing for diversity*. The Creative Independent. Retrieved November 18, 2022, from <https://thecreativeindependent.com/guides/how-to-begin-designing-for-diversity/>.

Rockley, A., Cooper, C., & Abel, S. (2015). *Intelligent content: A primer*. Laguna Hills, CA: XML Press.

Swisher, V. (2014). *Global content strategy: A primer*. Laguna Hills, CA: XML Press.

Swisher, V., & Preciado, R.L. (2021). *The personalization paradox*. Laguna Hills, CA: XML Press.

Van der Meij, H., & Carroll, J. (1995). Principles and heuristics for designing minimalist instruction. *Technical Communication*, 42(2), 243–261.

Verwijs, C. (2019, April 25). *Here's what's wrong with maturity models*. Scrum.org. Retrieved November 18, 2022, from <https://www.scrum.org/resources/blog/heres-whats-wrong-maturity-models>.

CHAPTER 5

Customer Experience and Content Operations

Kevin P. Nichols

Obsess about customer needs—and use that as your compass.

—Forrester, 2021

Do what you do so well that they will want to see it again and bring their friends.

—Walt Disney

Overview

Customer experience and Content Operations remain inextricably linked. When organizations fail to recognize this fact, their customer experiences result in fragmented or unproductive customer *content* experiences. Broken operations within a content life cycle yield subpar customer experiences—and conversely, successful ContentOps can yield more successful customer experiences. Today, omnichannel content delivery is crucial for many organizations. Omnichannel content experiences mean that the customer can start a task at one touchpoint (such as ordering a product online) and complete it at another (such as picking it up in a store) without feeling as if the experience is disrupted.

How to cite this book chapter:

Nichols, K.P. 2024. Customer Experience and Content Operations. In: Evia, C. (ed.) *Content Operations from Start to Scale: Perspectives from Industry Experts*. Pp. 75–90. Blacksburg: Virginia Tech Publishing. DOI: https://doi.org/10.21061/content_operations_evia_7. License: CC BY



Figure 5.1: Representation of a customer content omnichannel experience.
Credit: AvenueCX.

The concept of omnichannel delivery has emerged in the past few years to help organizations think through how to orchestrate seamless content experiences. Whereas multichannel delivery focuses on the processes and technologies for delivering content across an array of channels and devices, omnichannel focuses on coordinating activities across departmental silos to achieve a consistent message to users wherever they interact with an organization (OmnichannelX, n.d.). Figure 5.1 presents a graphic view of the omnichannel experience, with diverse channels of interaction, levels of participation, and stages through the process of a customer interacting with content.

COVID-19 presented the world—and businesses—with many challenges. For global brands, the pandemic forced workers to remain largely at home, and almost overnight, many businesses had to put much of their supply-chain management processes and content into cloud-based solutions. The importance of content within the customer experience became amplified, as customer interaction with content at every touchpoint was intensified by the demands of a shelter-in-place economy. Customer relationships with brands have always required content to support them, but never had the customer experience of content been so pivotal. During 2020, content was often the only element driving customer interactions with brands. As businesses struggled to maintain or retain their existing customer experience models while delivering their services and products, breakpoints and issues within their ContentOps workflows were exposed. Such issues could cause further fissures within already fragmented customer experiences.

This chapter explores the concepts of customer experience, customer journey, and content and examines their importance for brands and organizations. I will explore the relationships among these concepts before making the case

for where each falls within ContentOps. A solid understanding of the customer journey, and its stages, is necessary to frame its relationship to content and its role in ContentOps.

Customer Experience and Customer Journeys

A customer journey represents the relationship that a customer has with a brand, and it reflects the “path” that that customer takes while engaging with the brand. As a tool that brands leverage, several different types of customer journeys exist. This chapter explores a high-level customer journey map and specific customer journeys within it.

Let us turn for a moment to explore the relationship between customer experience, content, and brands. By “brands,” I mean how a person perceives an organization externally: the external manifestation of an organization within consumer perception. By “customer,” I mean any potential or existing obtainer or consumer of a particular brand’s services, offerings, or experiences. First, let’s start by defining a “customer journey,” a term familiar in marketing and customer experience departments.

Customer journey: Represents the broadest category of journeys and includes any and all types of customer journeys below. Simply put, a customer journey includes the “path” that a customer takes when engaging with a particular brand or organization. Often, “customer journey” is used interchangeably with “customer journey map,” but a customer journey map reflects a specific type of customer journey and contains specific parameters, which a journey—much more general—does not.

Customer journey map: This is the most common type of customer journey and the one with which most people are probably familiar. A customer journey map visually represents the life cycle of a customer’s engagement with a specific brand and includes the following elements:

- An actor (also known as a *customer type*, *customer profile*, *customer segment*, or *persona*). The archetype or profile upon which the journey is based.
- The end-to-end scenario. This represents the action taking place and often may refer to the entire life cycle that the map represents.
- Phases within the scenario. The stages, sometimes referred as “steps,” that reflect the phases of the customer’s path.
- Emotional or intellectual response. The thought processes or emotional response that a customer possesses as they complete each action or phase.
- Triggers. Any antecedents that trigger an action or propel the actor to take action.
- Opportunities. Noteworthy possibilities provided within each phase, which include potential content options, such as thought leadership insights or loyalty rewards content.

Brands use customer journeys to build customer-centric brand experiences. A well-executed customer journey provides the blueprint for who a customer is, how that person behaves, and what the customer needs. Thus, a customer journey can help a brand anticipate customer needs to create meaningful customer experiences: for the brand itself and its websites, digital experiences, in-store experiences, and content. Customer journeys enable brands to create experiences that reflect an understanding of customers and their needs. From the inception of the brand strategy through to the content to support it, customer journeys provide a tool to get everything that a brand needs to understand (correctly) about the customer.

Gartner conducts extensive research in customer experience return on investment (ROI) and recently concluded that customer-centric companies are 60% more profitable than companies that are not. A 2019 report found that “CX [customer experience] drives over two-thirds of customer loyalty, outperforming brand and price combined” (Gartner, 2019). The same report notes that to improve customer experience, “analyzing customers’ salient experiences and [customer] journeys often illuminates numerous suboptimal touchpoints that contribute to customers’ overall [negative] perceptions” (Gartner, 2019). Several recent research reports—including from McKinsey (Diebner et al., 2021), Forrester (and Sirius Decisions) (Forrester, 2021), *Harvard Business Review* (Gruner, 2021), and *Forbes* (Jenkins, 2021)—suggest the efficacy and necessity of customer journeys.

In order to put these concepts in context, Figure 5.2 shows a typical journey. You need not understand the intricacies of the map illustrated in Figure 5.2. It presents a simple customer journey map, of which a typical printout from a plotter machine can wrap around the walls of a corner office in a skyscraper building. But let us unpack the illustration for a moment. The stages (columns) present the “life cycle” of the customer—that is, their entire relationship with a particular brand. Each stage presents unique opportunities for the brand to build and/or evolve its relationship with the customer through delivering specific types of content:

- Awareness. When a potential customer first becomes informed about the brand.
- Consideration. When the potential customer ponders whether they should engage with the brand.
- Decisions (also known as conversion). When the potential customer becomes an actual customer.
- Delivery and use. When the customer receives the product or service and uses it; sometimes this phase bifurcates into another called “Support,” which refers to when the customer requires assistance for the product (needing to call customer service, for instance).
- Loyalty and advocacy. “Loyalty” refers to when the customer comes back to the brand to reengage, and “advocacy” refers to when the customer evangelizes the brand (for example, by sharing their experience on social media).

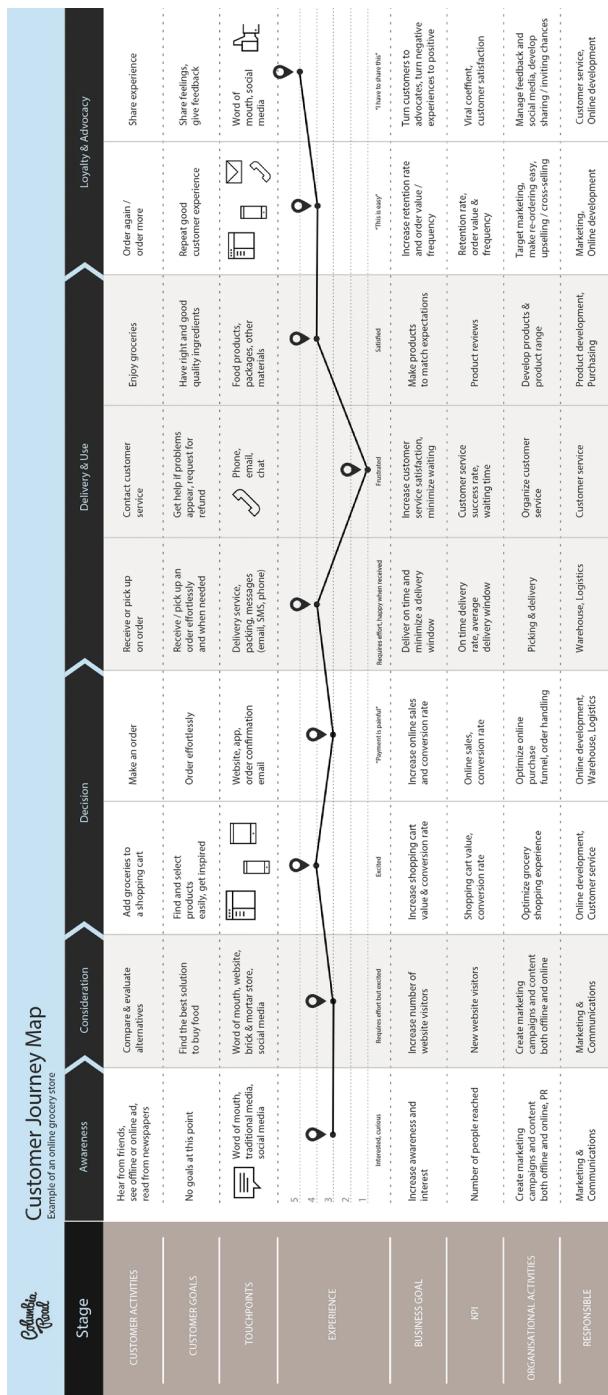


Figure 5.2: Sample customer journey map based on a fictional online grocery store. Credit: Columbia Road.

We often see these phases in a closed-loop diagram, where cross-selling and upselling, along with repurchasing, are factors that bridge the gap between loyalty and reengagement. In fact, McKinsey developed a model centered on the cyclical nature of customer interaction, called the “customer decision journey” model (Court et al., 2009).

The high-level paradigms presented in Figure 5.2 accommodate typical customer behavior. Most customer relationships with a brand have these high-level phases in some form or another. For this chapter, I will focus on the following stages: Perception (Awareness), Consideration, Conversion, Use, Loyalty, Advocacy, and Reengagement.

Each customer’s journey with a brand is different, and every customer has a unique journey in how they engage with a brand. Further, the activities that a customer completes in each of these phases, or “micro-journeys” within each phase, differ radically, depending on the industry type and customer type. That is why the second type of journey that we need to consider is a more specific customer journey within the larger map, based on prioritized tasks like “download a whitepaper,” “fill out a form,” “return a gift,” and so on. These more **task-based customer journeys** allow us to map content to specific customer tasks, touchpoints, and areas within a web page or printed page.

In customer experience, marketing and digital, we can segment certain types of customers based on behavioral conditions. To do so, we leverage quantitative and qualitative customer data and customer insights from analytics, consumer research, and user testing to create a picture of the customer profile or type. We create what we call **segmentation models** and **personae**. A persona represents an archetype of a customer based on shared behaviors. Without getting into too much detail, a “Soccer Parent” or “Digital Native” might represent typical personae used by a retail company to cluster a group of like-minded customers. We can then design journeys specifically for these customers and their behavioral profiles.

You might ask why we would do so, but consider for a moment the needs of a “Silver Surfer Adventurer” persona (defined as a retired seventy-year-old man who travels the world) purchasing different types of insurance versus those of a “Social Justice Advocate Digital Native” persona (defined as a thirty-year-old millennial trans person living in Denver, Colorado). And think about not only the products and services offered to them but how a US-based global insurance brand such as Safety Insurance would position its company, products, and services to those two customer types.

We design journeys specifically for personas or customer types. It is important to note here that we also create personae or profiles for our primary and secondary customer bases and for customers whom we would like to target (and may not be reaching). Generally, I recommend no more than five to seven personae or profiles so that a ContentOps team can limit these to key customer bases, which allows the team to prioritize decisions and think through strategy for the most significant audiences.

Content and Customer Journey

Now, let us return to those high-level phases, as these form the foundations of effective content messaging within customer experience—and, as a result, effective ContentOps workflows. We will explore them from a foundational perspective: what a brand must achieve to build its relationship with the customer within each phase.

- Perception (Awareness). Identifies the stage at which the prospective customer is **entirely unaware** of a brand and its service or offerings. Brands use advertising and marketing content to build initial awareness of their services or products for prospective customers (and to highlight new opportunities and cross-sell paths for existing customers). A brand has little time to gain the attention of prospective customers: it takes a customer “10 seconds or less” (Court et al., 2009) to form an impression of an organization through branded content. Therefore, *competitively differentiated, memorable*, and *effective marketing messaging*, and the most compelling “look and feel” to reach the *right* audiences, proves critical. (Think about that Super Bowl ad that you don’t forget the morning after the game.)
- Consideration. The brand’s content manages to gain the customer’s attention and the customer wishes to evaluate the brand prior to purchase. This phase includes any of the following: in-store appraisal, research on a company’s website, competitive analysis (comparing the company to its competitors), third-party peer reviews (Yelp, customer reports, and so on), and validation through influencers, social media, or known persons (friends/family). Product packaging, website content, social media content, and trained in-store staff (which means tight messaging, proof points, and talking points) remain crucial in this phase. Getting content right here is vitally important. Google’s own research confirms that “it takes .05 seconds for a visitor to form an opinion of a website” (Nunn, 2016). And the rule of seven typically applies to a customer prior to making a purchase, meaning that for B2B (business-to-business) and B2C (business to consumer) transactions, “It takes an average of seven interactions with your brand before a purchase will take place” (Cairns, 2021, p. 150). This phase also carries implications for technical communication teams, as reports show that 53% of global customers use technical product information to learn more about a product before purchasing it (Cision, 2017).
- Conversion. The customer buys the product. Here, the point-of-sale (POS) interaction proves paramount, as does the engagement with the sales representative and/or the delivery or courier service. Whether the delivery happens in one channel or in a multi- or omnichannel approach, the customer expects a seamless experience. Thus, if the customer orders the product online, they assume that they will be able to pick it up in store without issue.

All messaging and content supporting the conversion process remains critical, because the customer provides money for the product and expects a certain level of service and guarantee. In many industries, customers trust that their privacy will be preserved, and they certainly expect that their financial information will be protected. From the point of sale through the delivery of the product, all communication should evoke trust, integrity, and security. All content and branded information should communicate this fact. Any breakdown here can create an irreparable fissure between customer and brand. Processes should be streamlined and optimized—you want the customer to convert as efficiently as possible—and they should observe contemporary trends in showing respect to customers, such as offering contactless purchases in light of COVID-19 safety protocols.

- Use. This phase constitutes the evaluation stage, or trial period, of a product, where the “proof is in the pudding.” This phase becomes the nexus that solidifies (or breaks) customer trust. True, the product must stand the test of how the brand previously positioned it in the customer’s experience through its marketing and advertising content. Authenticity and true representation of the product and its capabilities remain essential in effective brand building. But the product’s packaging, instructions, any content on the website meant to troubleshoot the use of the product, and/or the training for the in-store salesperson (or sales representative who may receive a call) must all support a positive customer experience in using the product. Here, the brand should include consistent messaging on how to use the product and anticipate any customer needs regarding the product’s use. This process also highlights the need for robust customer journey modeling at a more micro level to ascertain customer needs. Smart brands develop content to support the product and its use for the customer—and the people and technologies (e.g., chatbots) required to support that product’s use—to ensure an exceptional customer experience. For medical products, clear and concise labeling that evokes safety, health, and care remains an essential component in ensuring customer trust. And even *smarter* brands anticipate that the customer will appreciate the product and create entry points for the next phase, loyalty and advocacy: loyalty through incentives to retain the customer, and advocacy through campaigns to inspire the customer to share their experience with the product.
- Loyalty and Advocacy. I combine these two phases, but in actuality these denominators remain distinct, with loyalty pertaining to customer retention and advocacy to customer influence. Loyalty is the ultimate goal of successful brand building, and advocacy is the cherry on top. Loyalty means that the customer remains faithful (although brands should not fall into the trap of hubris, as arrogance may lead to customer experience failures, yielding customer attrition). The cost of losing a customer proves detrimental to brands and has been quantified by numerous studies. Customer deflection

metrics also contrast the cost of customer retention with that of acquiring new customers. There is no comparison: the general consensus is that for most industries, it costs five times as much to acquire as opposed to retain customers. Loyalty comes from satisfying customers with good customer experience, consistent brand messaging, reward programs (if these exist), and loyalty programs (if these exist). Advocacy derives from incentives to attract influencers, such as campaigns that encourage the customer to share their experiences with the product and make it easy for them to do so. Interestingly, customers place the highest trust in referrals on brands from people whom they know, and word of mouth marketing “drives \$6 trillion of annual consumer spending, accounts for 13% of consumer sales, and people are 90% more likely to trust and buy from a brand recommended by a friend” (Qualtrics, n.d.). Here is where positive customer experience at POS and in-store experiences provides critical touchpoints, as these customers remember human interactions more than they do digital ones.

- Reengage. Many brands break down in this phase, but smart ones don’t. The primary concern here revolves around replicating the triggers that drove the customer to purchase. An additional task is appealing to any brand loyalty that exists from previous engagement.

At this point, you should have a basic understanding of the customer journey and its role in customer experience. So what does this mean for content and Content Operations? Customer journeys allow brands to consider each stage of their customers’ interactions—brands typically create a few journeys for five to seven personae—to ensure that the brand has the content and messaging that speak to customers’ needs in each journey stage. Brands can leverage each stage in developing every aspect of their identity and messaging in a customer-centric manner for brand identity content (logo, visual look and feel), marketing messaging, content that positions and supports the products, and the like.

This process also includes programs that the brand develops to meet customers’ needs throughout all of its service-delivering channels. Customer journeys can help identify customer touchpoints and help brands determine what content or interactions customers need in which channels (for instance, web, mobile, or in-store). Because customer journey maps consider the entire customer life cycle, they help to establish a way for content to meet the customer throughout that life cycle. Thus, the brand considers every stage of the journey when building out its identity in order to meet the customer where they are and where they need to go.

Micro-journeys, or task-based journeys, help a brand figure out content across channels and across more micro-steps. These journeys can help brands identify more specific content. We use content journey maps to identify content at the **content type** level. In “Content Type,” Jonathon Colman defines a content type as “a specification for a structured, standardized, reusable, and mutually exclusive kind of information entity” (2014, p. 48). Examples include thought

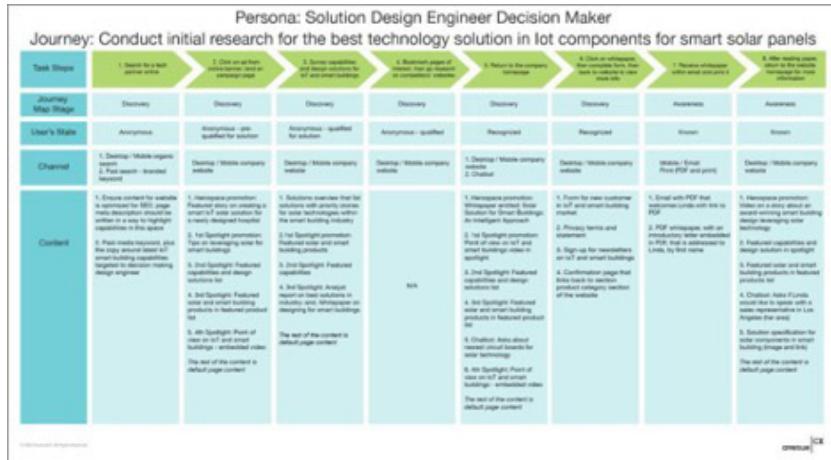


Figure 5.3: Task-based customer journey with content mapping. Credit: AvenueCX.

leadership white papers, event pages (for a conference), a website home page, and a frequently asked questions (FAQ) page. A task-based customer journey allows an organization to map content at a more structured level; we can actually map content to specific elements within a content type—for example, to a hero space within a website. Figure 5.3 shows an example.

In Figure 5.3’s example, you can see how content is mapped to specific areas within a website. This exercise illustrates the need for a structured content solution to support such a model. Within content management, structured content design is a method for architecting and conceiving content.

This concept is critical for effective content management. Best practices for structured content design always start with defining a core set of content types that the content experience requires. Defining content types results in requirements and specifications for all content entities within a content management system (CMS). The content type defines the structure for a template, page type, or container of content. The content type should describe the intent of the information. By defining content based on its type, an organization can obtain several benefits, including the following:

- Improve user experience (UX) to help users find the content that they need more quickly through consistent categorization and search functionality.
- Deliver consistency of structure, message, voice, and style within content type.
- Align content types to products at introduction and customer journey stage.
- Reduce/eliminate time spent doing “governing” work on what content types are used (and when, where, and how).

- Provide guidance for content creators with specific content types and structure, thus eliminating ambiguity.
- Lay the foundation for advanced UX, utilizing machine learning/artificial intelligence, A/B versioning, personalization, and omnichannel interaction to support business goals, such as increased cross-sell/upsell or other marketing and sales goals.
- Create clean user analytics that can identify content areas to focus on, expand, or remove as well as highlight investment areas.

All these benefits prove relevant when considering the customer journey. Once a content type is defined, it becomes the basis for determining which templates or structures are required to support it. A structure comprises the building blocks of a content type, and a content type template will break down into modules that may or may not be called components. Some content management systems brand this terminology differently, but think of it this way: a home page is a content type, which could also be a template. Within a website's home page, you might have hero space (component or module), featured products (component), news and events (component), a featured story (component), and other content.

Now let's get back to that customer journey. As you can see in Figure 5.3, a task-based customer journey allows you to map content to a journey. Thus, having a solution in place for creation, publication, and delivery of that content remains critical if one wishes to support an optimized customer experience. This means placing the customer journey as a primary requirement within ContentOps and, equally important, technical stack requirements. To do this, an organization must have a customer journey in place, and it should also have a plan to map content for priority tasks (micro-journeys) within it—or already have done so.

Considerations for Omnichannel within Customer Journeys

For omnichannel experiences, customer journeys become even more critical—and present even more complexity with content solutions—because of the amount of integration involved with strategy, content, and technologies. If we consider even a typical number of channels within the omnichannel customer experience, we can see why: websites, mobile apps, kiosks, in-store interactions, POS interactions, customer sales representatives, customer support representatives, print (e.g., packaging, manuals, brochures, or correspondence), events, radio, television, and digital signage. An omnichannel approach includes a business strategy that considers the customer in all of the brand's touchpoints and seeks to deliver content and services to that customer regardless of where they are within a journey.

Omnichannel strategy includes an approach to deliver content to several of the above channels, but it optimizes that content to what the customer needs within that channel based on their interactions. Some key concepts that

omnichannel emphasizes include a single view of the customer, or the understanding of that customer regardless of which touchpoint is leveraged, and integrated product inventory, which allows for a product to be ordered online and picked up in store. It also includes transparency in ordering and procurement for the customer.

Insofar as content is concerned, an omnichannel approach requires cross-channel analytics collection and cross-channel content publishing. This integration may include integrating an email system, website, and customer support center so that the organization understands what a customer has done and is able to respond to their actions in real time. A robust omnichannel model requires significant investment by an organization as well as an enterprise approach to data and content management. It also requires advanced personalization capabilities to tailor very specific content to a customer based on much more than just what content they have viewed. This personalization should consider all stages of the customer journey, especially once the organization knows who that customer is. Although most businesses will not adopt a true omnichannel strategy (and they may not necessarily require one), many can take away lessons from an omnichannel mindset, such as optimizing customer experience across the various stages of the customer journey by ensuring that there is enough content in each channel to support the needs and use cases necessitated by those touchpoints.

Personalization and Customer Journeys

Personalization refers to any content contextually targeted based on the following:

- who someone is,
- where someone is,
- what someone does,
- when someone does it,
- how someone does it, or
- why someone does it.

Within this context, localized content is technically personalized content, because it is contextual to a user's location (or language preference, regardless of location). Although many businesses leverage personalization solely after a user authenticates and then personalize content against the user's profile or previous purchase patterns, there is also much benefit to other forms of personalization. For example, if someone shows an interest in a particular industry vertical or product, showcasing content relevant to that topic can offer more tailored content to what the user seeks—even if that user remains unknown or unauthenticated. Other tactics include the following:

- Personalization based on key terms searched (if you know the user is interested in a particular topic).
- Cross-selling and upselling based on product views or previously purchased products.
- Recommended content based on the click-stream of content topics viewed.
- Content based on the company or organization where the user is located—triggered by geolocation based on Internet Protocol (IP) address.
- Specific promotions based on geographical location or region.
- Targeted content based on paid media that the user clicked on or an email message that they opened to get to the website. (Note: advanced personalization can also change the content in email prior to the user's opening it, based on user behavior on a website or other channels.)
- Content that pushes the customer from one stage of the journey to the next.

Today, most major brands offer some form of personalization on the consumer and B2B websites. Done well, it can yield all sorts of benefits. McKinsey reports on these (Lindecrantz et al., 2020). This report notes that 80% of customers want personalization from retailers (but only if it is done well and provides relevant experiences) and that it can yield 20% higher customer satisfaction rates. The report provides many more useful statistics as well as key considerations when thinking about first-time business personalization. For B2B considerations, the *2020 B2B Buyer Behavior Study* from Demand Gen (Demand Gen Report, n.d.) offers some of the following conclusions:

- 76% of those surveyed expect personalized content specific to their needs.
- 70% stated that it was very important to view “relevant content that speaks directly to our company.”
- 96% wanted content that “spoke directly to our industry needs.”

Study after study has quantified the value of personalization for businesses, whether through an Account-Based Marketing model for the B2B or a targeted campaign for a retail customer. It can also be an important part of a business strategy to acquire new customers and drive efforts such as gated content strategies and lead qualification for prospective customers. Customer journeys play a primary role in driving effective personalized content experiences.

Content Operations and Customer Journeys

All of the above indicates that customer journeys present many key areas of consideration for ContentOps. A customer journey should inform requirements in the design of the following areas of ContentOps and operations models.

- The enterprise content strategy/platform-specific content strategies. Use customer journeys as a key input to inform strategic decisions and priorities around content objectives, internal operations objectives, and performance objectives.
- Content life cycles and workflows. Look to design workflows that will support the various aspects of cross-channel customer journeys. Use customer journey maps to define the workflow at the content type level.
- Content structure. Use task-based customer journey maps as a primary input into the design decisions and rules for how content modules or components are used or reused across platforms.
- Content governance. Leverage customer journeys as a process and primary input within the governance model; ensure representation within the governance structure of the journey.
- Content technology requirements. Use both customer maps and task-based customer journeys to define detailed requirements for technology stacks. Customer journey maps can identify points of system integration, as can cross-channel task-based maps.

The above considerations provide a blueprint for how and where customer journeys plug into ContentOps modeling and thinking.

Content’s Role in Quantifying Customer Experience

Customer journeys can help identify which content to consider in a given transaction. This puts content in the role of having the ability to improve or move the needle on customer experience. An optimized ContentOps team and model, one that demonstrates improved content, can then improve the customer experience. These points are important; a ContentOps team should think about how it can implement ways to measure this impact. Content remains a valuable tool to quantify customer experience, which adds to its value as a business asset. A ContentOps workflow should look at the customer journey with the customer experience team and define a metrics approach that demonstrates the role that content plays in customer experience performance. By demonstrating this impact—and then cross-referencing these metrics with internal operations metrics, such as cost to produce content—a ContentOps team can demonstrate how content proves its worth as an asset within customer experience.

Conclusion

This chapter has made the case for customer journeys, demonstrated the components that make up customer journeys, and explained specific use

cases. A solid understanding of each of these areas is critical in understanding the roles that customer journeys play in regard to content. This understanding provides the foundations for recognizing where customer journeys plug into ContentOps. Effectively leveraging customer journeys will position a ContentOps model for success in any organization. Tying ContentOps to customer experience strategically can also help elevate the role and importance of ContentOps within organizations.

Customer journeys require complicated content solutions. A solid and firm operating model for content production—and ContentOps team—is essential for the optimal performance of a journey.

References

Cairns, A. (2021). B2B marketing strategy: Finding the needle in the haystack. In U.G. Seebacher (Ed.), *B2B marketing: A guidebook for the classroom to the boardroom* (pp. 121–156). Cham, Switzerland: Springer.

Cision. (2017, December 7). *SDL research signals shift in customer buying behavior*. Cision PR Newswire. Retrieved June 17, 2022, from <https://www.prnewswire.com/news-releases/sdl-research-signals-shift-in-customer-buying-behavior-662532603.html>.

Colman, J. (2014). Content type. In S. Abel & R.A. Bailie (Eds.), *The language of content strategy* (pp. 48–49). Laguna Hills, CA: XML Press.

Court, D., Elzinga, D., Mulder, S., & Vettvik, O.J. (2009, June 1). *The consumer decision journey*. McKinsey & Company. Retrieved June 17, 2022, from <https://www.mckinsey.com/business-functions/growth-marketing-and-sales/our-insights/the-consumer-decision-journey>.

Demand Gen Report. (n.d.). *2020 B2B buyer behavior study*. Retrieved June 17, 2022, from <https://www.demandgenreport.com/resources/reports/2020-b2b-buyer-behavior-study/purchase-plans-still-progressing-despite-disruption-but-with-increased-expectations-for-relevance-personalization/>.

Diebner, R., Malfara, D., Neher, K., Thompson, M., & Vancauwenberghe, M. (2021, February 24). Prediction: The future of CX. *McKinsey Quarterly*. Retrieved June 17, 2022, from <https://www.mckinsey.com/business-functions/growth-marketing-and-sales/our-insights/prediction-the-future-of-cx>.

Forrester. (2021). *Planning assumptions for B2B 2022*. Retrieved June 17, 2022, from <https://www.forrester.com/planning-assumptions/b2b-marketing-guide/>.

Gartner. (2019, January 16). *Creating a high-impact customer experience strategy*. Gartner Research. Retrieved June 17, 2022, from <https://www.gartner.com/en/documents/3899777>.

Gruner, R.L. (2021, May 14). 4 strategies to simplify the customer journey. *Harvard Business Review*. Retrieved June 17, 2022, from <https://hbr.org/2021/05/4-strategies-to-simplify-the-customer-journey>.

Jenkins, S. (2021, February 2). *Customer journey mapping and its essential place for today's modern marketer*. Forbes. Retrieved June 17, 2022, from <https://www.forbes.com/sites/forbesagencycouncil/2021/02/02/customer-journey-mapping-and-its-essential-place-for-todays-modern-marketer>.

Lindecrantz, E., Tjon Pian Gi, M., & Zerbi, S. (2020, April 28). *Personalizing the customer experience: Driving differentiation in retail*. McKinsey & Company. Retrieved June 17, 2022, from <https://www.mckinsey.com/industries/retail/our-insights/personalizing-the-customer-experience-driving-differentiation-in-retail>.

Nunn, C. (2016, February 11). 11 branding stats crazier than Miley. *ActionCard*. Retrieved June 17, 2022, from <https://actioncardapp.com/2016/02/11/11-branding-stats-crazier-than-miley/>.

OmnichannelX. (n.d.). *What is omnichannel?* Retrieved June 17, 2022, from <https://omnichannelx.digital/what-is-omnichannel-the-omnichannelx-conference/>.

Qualtrics. (n.d.). *The ultimate guide to customer loyalty*. Retrieved June 17, 2022, from <https://www.qualtrics.com/experience-management/customer/customer-loyalty/>.

CHAPTER 6

Personalization and Content Operations

Jeffrey MacIntyre

Personalized experience is a uniquely vexing topic in technology and design circles. It's simultaneously a paradisal and dystopian scenario among consumers—and a variously utopian or fraught one among practitioners who manage them. Who could forget the Amazon chestnut that 35% of total revenues derive from its recommendations feature set, or the moment when Netflix navigation and browse features went fully algorithmic (bye bye, default state!), or when Spotify's Discover Weekly playlists started to make waves?

Fewer know that Spotify did this quietly, on the back of a metadata engine acquisition, to "smarten" how it curated its content catalog. "Information Architecture (IA) before Artificial Intelligence (AI)" is more than a catchy slogan: it aptly describes the role of structured content and data in the destiny of every connected experience. Amid the polarized perceptions that it is either magical or malicious, designing for personalization is also the source of considerable reward in business: the results, when realized, can be sizable, and the attribution to content's performance and design decisions quite clear. The enduring arguments for a connected experience run something as presented in Table 6.1.

Sounds dreamy, right?

There is a personalization gap that clouds the clarity that we would expect, however, twenty-plus years into the history of the web, and it's the sum of all fears and fantasies that this technology tends to evoke, particularly for those who see it as more the province of high-powered, digital-by-default organizations than

How to cite this book chapter:

MacIntyre, J. 2024. Personalization and Content Operations. In: Evia, C. (ed.) *Content Operations from Start to Scale: Perspectives from Industry Experts*. Pp. 91–108. Blacksburg: Virginia Tech Publishing. DOI: https://doi.org/10.21061/content_operations_evia_8. License: CC BY

Table 6.1: The four enduring arguments for personalization.

	Business-centric	Customer-centric
Front stage: Your customer experience (CX)	Efficiency (optimization and automation)	Experience (a dynamic user experience and connected experience-powered features)
Backstage: How you operate	Orchestration (internal coordination of systems and data)	Understanding (deepened customer intimacy)

ones like their own. Put simply, the refrain goes something like this: “We’re not Amazon, and we’ll never be.” Like any pernicious framing, it’s not entirely untrue. But the personalization gap has disadvantaged many from succeeding in their journey to increased customer-centricity.

Simplified customer experiences dominate in businesses of every type as a leading source of aspirational competitive advantage. Scale and systematization are key to the thrust of ContentOps, so questions of automation or personalization are never far off. ContentOps has both a role and a reward in significantly advancing the state of practice here—and in time, folks may rightly see connected experiences as the natural fruit of a strong Content Ops function.

The ContentOps Advantage

The essential riddle that has bewitched countless product, content, and information professionals pursuing how to design and deliver personalized experiences can be summarized by a quip often attributed to children’s author Theodor Geisel: “Sometimes the questions are complicated and the answers are simple.”

The right content to the right person at the right time is the end goal that we all know, and it’s a commonplace for content professionals—even those operating in the default state of a user interface, absent any personalized tweaks. Yet as any digital professional worth their salt also knows, delivering on “simple” is a remarkable skill and an all-too-rare feat achieved in our projects. Accenture calls personalization the challenge of how “to uniquely serve everyone without overwhelming anyone” (Accenture Interactive, 2018, p. 2) in its customary setting: a digital interface. This represents a massive reverse-zoom from the experience of humans navigating 50,000-square-foot shopping experiences with 100,000 products merchandised physically to a 4-inch black rectangular screen regaling us with 10 million items to be digitally discovered.

Getting to “simple” for users is anything but. The trick is working with the grammar of personalization—the who, what, when, where, and how

parameters that fire every given interaction that you personalize, customize, or automate—to bend the complex to the clear and confidently narrow end users' choices from bewildering to cohesive. The complication lies in wrestling content, audience, and context into a meaningfully generalizable set of entities that can work satisfactorily and consistently.

The good news? The complexity quarterback we're talking about sounds a lot like you.

Designing for personalization rewards a close, root understanding of one's content offering, life cycle, and mechanics. Systems whisperers are just the right kind of folks for this simplification request. User interactions depend upon the coordination, if not outright orchestration, of multiple systems of record: content management system (CMS), digital asset management (DAM), customer data platform (CDP), product information management (PIM), and marketing automation platform (MAP), to name a few of the familiar tech-stack suspects. All connected experiences demand an IA-before-AI foundation to ensure the reliable flow and understanding of data and information. In capable hands, ContentOps is the stewardship of this information layer.

A ContentOps outlook provides precisely the preconditions, sensibility, and tools necessary for closing the personalization gap and realizing these simplified experiences as more likely outcomes in your own projects. For this reason, we fully expect in coming years to see more ContentOps practitioners recognized as leading lights of the nascent personalization field of practice.

Closing the Personalization Gap

Step zero in achieving personalization success requires understanding the outsized role that failure plays in our understanding of what has not worked. For years now, it has been customary to see the fortunes of personalization wax and wane, going from glitter to gauche. Many product pros, particularly in Silicon Valley, sing its praises; just as many design folks, especially strategists and content management wonks, are quick to call out personalization as an in-practice disappointment.

For content marketing consultant Robert Rose, personalization is “the sexiest thing no one is doing yet.” Its effects are generally approved by consumers, results are roundly appreciated by organizations, its rank as a business priority remains high—and for all that, the mechanics of execution remain a contentious issue in practice.

Information management professionals, from the dawn of digital to today, have said that single sourcing is the path to liberating a content offering for personalized, customized, or automated expression. If we create well-defined atomic units of content, we can trace and maximize their reuse. This somewhat utopian project confounds many efforts to scale personalization.

Before we explore why the Edenic model and its totalizing logic—“it’s not personalization unless we personalize all the things”—undermines such efforts, consider also the consumer and business backdrop.

For the longest time, software vendors have steadily touted the seemingly ongoing arrival of personalization at scale. Vendors in the CMS, MAP, and personalization engine categories have been relentless in expressing connected experience as the next logical evolution in experience design. At the same time, different digital professional focus areas from content, to data, to product have taken varied approaches in their contributions to personalized customer experience (CX). Content subject matter experts (SMEs) are wont to talk about componentization, the modeling of content assets into modular chunks. Data SMEs may be more inclined to involve testing and experimentation generally in the world of experience optimization. Product SMEs, by comparison, may take a more functional-level and feature-engineering approach, providing particular tactical automations, customizations, or personalizations in their products. The thinking here has been divergent for years, and because the software market is similarly fragmentary and aimed toward these functional audiences, it’s unsurprising that a holistic playbook for broad-based personalization has eluded us.

Who are the poster children for connected experience? It is indeed easy to gesture at MAANA companies (Meta, Alphabet, Amazon, Netflix, Apple) and proclaim that we are a full generation into users who are broadly habituated to the algorithmic age of user experience. It is similarly easy to point out the many shortcomings of connected experiences, especially in black-box machine learning techniques that generate AI that is unexplainable, and unaccountable, to civil society. Long-standing issues of overfitting and overdetermined targeting dog our day-to-day interactions online, and stories of “perso-fails” firing off overzealous notifications are a kind of ongoing punchline. A growing awareness of the distortions of algorithmically defined settings and interaction patterns dominates the news.

Personalization is so synonymous with both prize and peril in digital circles that it has become in recent years a reliable marker of deeply unfashionable, starry-eyed tech optimism. Personalization, in this context, is invoked more as a joke about green practitioners who have yet to be roughed up by a tour of duty on their own failed perso project with a gaudy technology that cannot deliver even limited returns on the good-as-DOA pledge of the right content to the right person at the right time.

The cognoscenti aren’t generally wrong about the hazards of personalization efforts, but they’re often hostage to a glib disregard for significant advances in technology and business practice that mean fewer people make the same mistakes any longer. For content subject matter experts, in particular, there is a sunny future ahead and it starts with the tool kit that they bring to the table.

Why haven’t we had this conversation to properly size up this opportunity?

Personalization is a broader mosaic of various expressions of digital interactions: websites, apps, notifications, product and content recommenders, multichannel communications like email-based nurture campaigns or SMS-sent automations, conversational user interfaces, and complex sets of rules and algorithms—and data management tools, too. Across all of this, we've lacked the ability to develop a framework to design holistically for personalized experience. What pockets of SME practice do exist can largely be found in the domains of specific software categories, like content management, taxonomy and metadata management, analytics, experience optimization, content and product recommendations, and marketing automation.

What's more, the draw to personalization has become better appreciated in the broader context of product differentiation, information overload, and increasingly clear data about consumer preferences for simplified, streamlined customer experiences. For example, the analyst firm Gartner managed to set off a tempest in a teacup for the technology-obsessive when it recently proclaimed that some 80% of organizations aspiring to personalized customer experiences will falter and eventually abandon them altogether (Gartner, 2019). A fine scaremongering tactic, but it does helpfully remind us that a personalization gap remains when it comes to our professional acumen and our practices. Meanwhile, user experience (UX) practitioners have been particularly lagging behind their peers in marketing and analytics as to the art of the possible. The design of personalized experiences and the advance of a conversation about design's role in personalization effectiveness have likely suffered as a result.

Gartner didn't come to bury personalization: it came to pronounce magical thinking dead. "Big bang" or moonshot personalization doesn't work in any organization, no matter how personalization-curious it may be. It is simply impossible in the face of existing experience and metadata debt common to most companies. I'm fond of a question a consultancy in this space once used to gauge the sprawling power of personalization and optimization software: If you can do anything (and you can), how then do you decide to focus your time and efforts?

The quintessential problem of personalization is that the software so surpasses the available talent and experience in wielding it effectively and efficiently. Meanwhile, personalized features and functionality have sneak-thieved their way into all kinds of software and design practices, with novel digital interfaces replacing their customary human originals. Designing for them may be exotic to you, but connected experiences are fast encroaching into the lives of everyday people. Consider the prosaic fast-food restaurant visit: at a bare minimum, the McDonald's restaurant connected experience spans IBM Watson AI voice interface at the drive-through, a Dynamic Yield-powered dynamic menu and kiosk system, and a loyalty software responsible for doling out appropriate deals both pre- and post-order and, conceivably, pricing in its mobile app.

The personalization gap is not about technological shortcomings. It is not about talent, though it is (anecdotally speaking) unmistakably affected by relative inexperience among many people confronted with personalization projects. The personalization gap is the unserious conversation among digital professionals that misses the critical contribution that design offers, and that students of ContentOps can utilize, as a remedy to all this confounding complexity.

Progressive Personalization, a Playbook for Connected Experience

The world is noisy. Navigating it is tiring at times—and uniquely taxing in the information-dense world of digital products and services. The UX world has long maintained that better design and technology create a balm for the digital disorder in our daily lives, with personalization heralding a glorious advance: a way of radically simplifying and calming our user interfaces and screens while improving their effectiveness.

We have been doing personalization wrong, particularly in the crusading focus on totalizing, moonshot-or-bust end states that are rarely practical, efficient, or sustainable. Given the success of the MAANA giants, it's easy to paint personalization approaches as out of touch with the Main Street business world or the needs of governments or nonprofits. The real-world line on personalization is that it requires too much and that risks exceed rewards. Enthusiasts of personalization technology and applications come off as impractical at best, and at worst, out of touch.

Maybe both parties are exactly right.

Personalization needs a practical, pragmatic basis in order to be successful in business settings where we haven't even mastered our own content management and marketing technologies. And the fanciest design system is coming up short if it can't create a vernacular for how content is being contextually delivered based on different heuristics and semantics.

In short, getting to the essence matters. That takes a few forms.

1. IA before AI is axiomatic. The best algorithms arise from a foundational architecture of data and content, and personalization has a clear grammar. (The clearest “perso-fails” erupt when IA is not in the loop, as we'll see later.) Understand the map of the constituent components of your digital footprint across channels, like web and email and kiosks, and across end points, like websites and apps. Understand how content and data flow across them. Content must be componentized—but don't go overboard. There are on-ramps that can train your organization in what level of componentization is ripe for strategic opportunity.
2. A jobs-to-be-done or top-tasks paradigm will help bring added focus to one, or maybe two at most, goals for a given screen. How will a user

progress in an interaction flow that is itself seated in a wider customer journey? The definition and fidelity of these understandings are crucial, illuminating what gets personalized and when.

3. A culture of experimentation is unquestionably helpful. Testing and optimization are fundamental to successful personalization programs. They are also proving grounds for teams able to make rapid, ongoing changes to core content and any given interactions in the digital footprint. Nimble teams win at so much in digital, even scale plays.

There are whole pre-personalization projects that some organizations undertake here to ensure some base readiness. Many never evolve beyond them. Truthfully, their value to an organization that, say, implements a personalization engine is that the above work can guide an approach to using the software's features and functionality to simplify the customer experience (as opposed to only complicating your administration of it). The signature feature of every personalization software suite is its sheer potency—the innumerable interventions possible.

But there is a root grammar to all personalizations and automations and customizations that is good to remember: a what, a who, a when, a where, a how, and ... yes, a why. User interfaces usually state these as strings: "I want to do X, at the time of Y, for Z users, using these specific content elements."

These parameters form the basis for everything done, and so it makes sense that the components of your IA form the direct basis of your "where" (e.g., placements of a component on a given page template). The "who" are audiences that you already know and track. The "when" maps to discrete moments or interactions patterns in user flows, nested within a customer journey. The "how" will vary based on one's setting, but, broadly speaking, it includes dynamic content elements (pop-ups, in-line messages, ribbons, call-outs, or outright component swap-outs), recommenders, automations and campaigns, notifications, conversational user interfaces, and even customizations (i.e., user preference settings).

Consider how Slack messaging software instructs itself on whether you want to be notified of something in their product (Figure 6.1).

Sound overwhelming? In the absence of guiding artifacts, it absolutely is. Even with them, it can still be more than a little paralyzing. These are consequential decisions.

Designing for contextual content delivery is unlike designing for the default state, and so it comes with additional overhead and practical consequences for experienced owners. Hick's Law holds that a user's decision-making time lags in proportion to the number of choices presented to that user (Soegaard, 2021). Adherence to a clear and consistent vernacular in personalization practices goes a long way toward more stable and preferable outcomes. The more generalizable the building blocks of your customer experience, the easier it becomes to mold and transform it in a way that you can measure and appreciate its performance with end users—and for your organizational leadership.

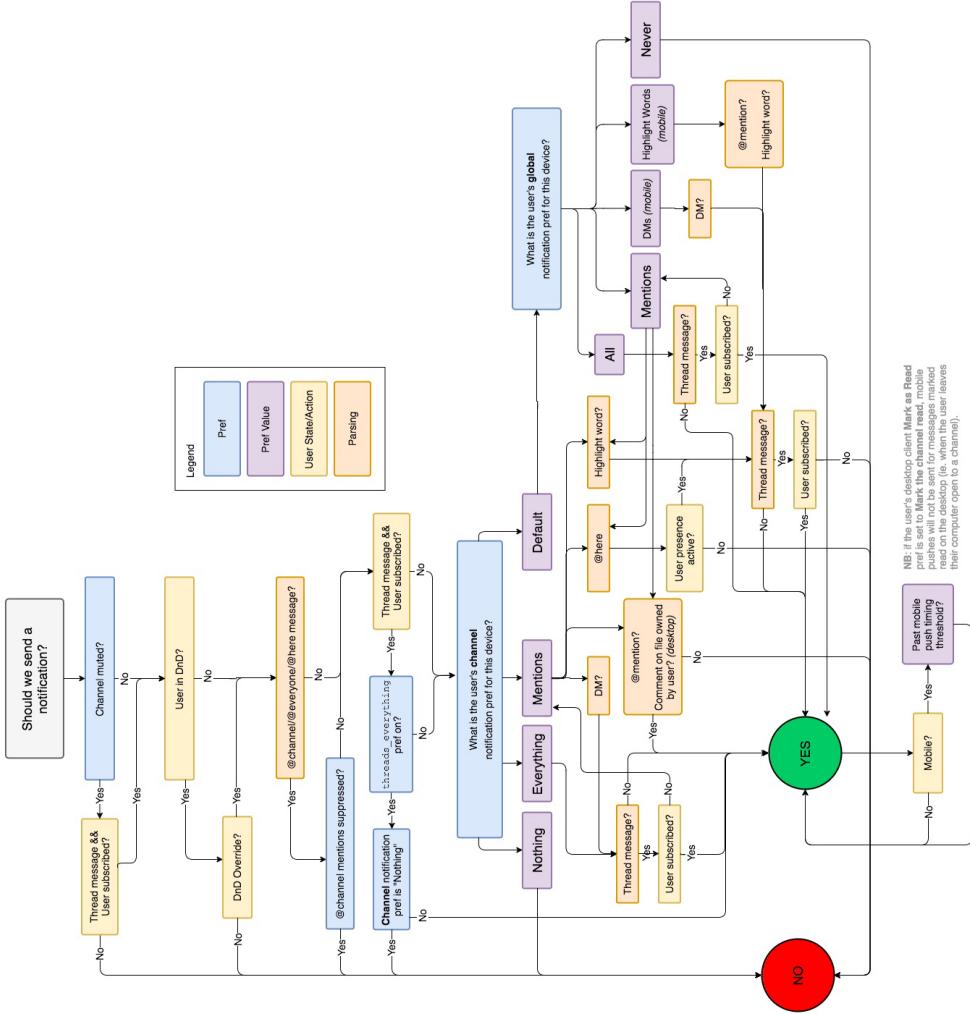


Figure 6.1: A framework that specifies the mechanics behind Slack notifications. Credit: Slack Engineering.

If you experience your own choice fatigue about what automations to run, say, it may be a fair indicator that the building blocks you work with are too detailed and specific. Try moving up the ladder to a higher-order generalization. Too many audience segments to juggle? Find parent categories. Too many page components in play? Only target header components. Certainly, avoid creating new classes and segments and interaction patterns that are not already part of how you manage and express your digital experience. Designing for connected experiences should be about selectivity, although it becomes important to be able to define for breadth in creating effects at levels that become statistically significant. If you are having trouble measuring the impact of changes to a single page, reconsider the number of pages in which you may be able to apply the tactic to reach the goal metric sooner.

That said, progressive personalization is a perspective of getting to early value in a connected experience, and sometimes the guiding light to sponsoring, say, a more in-depth customer journey assessment, or a “just-enough-taxonomy” metadata enrichment activity, or selecting and implementing a marketing automation platform, is itself a proto-personalization project. In progressive personalization, there are three classic on-ramps to connected experience, and success in one or more of them typically drives the sponsorship of increased investment and rigor in that connected experience. These on-ramps are only sequential in terms of relative time and resource investment, but they are standalone projects: the third on-ramp does not require completing the first and second, for example. They also encourage learning the ropes of one of three different experience modes of a connected experience: automation, customization, and personalization. This is noteworthy, since one mode may be wholly appropriate and others not, depending on the type of “smarts” that an organization is seeking to bring to their CX, or the type of hypothesis being tested. From the musical-motif names I’ve given them, it should be clear that they entail an increased level of committed initial and even ongoing effort—and that they involve increasing levels of sophistication around orchestration of content and data, too.

The Solo

This project finds you automating a key moment in the customer journey, systematizing it as an email-based set of interactions. This can be pulled off with a variety of freemium to enterprise email software, effectively requiring the ability to create an email drip campaign or series of automations. If engagement via this channel performs well, consider this a signal to expand utilization to other key moments to be delivered via the push mechanics of email or messaging, as appropriate. In terms of key performance indicators (KPIs), this is intended to drive journey progression, increase marketing qualified leads (MQLs), reduce churn, and increase reach.

In one instance, a software as a service (SaaS) client took initial steps to building out their marketing automation through this on-ramp. They identified a few select moments in their customer journey (a welcome flow, a nurture flow, and a “wingback” flow), creating email campaigns to automate those activities that were otherwise actions only available on their website or (at times) web app. The operative shift here was that this was a push mechanism rather than one that depended on continued user engagement and interest. In addition, the campaigns were automated as statistically relevant results confirmed their effectiveness. This implementation created a foundation for their future marketing team to run with—and paid off the entire automation project budget by setting up just one automation, to circle back via email with those whose credit card purchases were somehow interrupted.

The Sonata

This is the least effortful of the on-ramps from execution and technology standpoints. The CX specialist will clone a page (or stack of grouped pages) and then de- and reconstruct their content to appeal to a narrower audience. This could be as simple as copying an existing landing page and then rewriting and styling it for a specific customer segment. The notion here is to explore in a split test how the two (the as-is and the customized page[s]) perform. Does targeted content and messaging improve conversion rate among the target audience? This approach can be undertaken in the context of a broader Account-Based Marketing (ABM) effort or any other customer-centric movement in an organization. It will be important to make search engine optimization accommodations for this approach, and similarly, one will need to be practical about what this would take if scaled up as a segmented experience beyond a single cluster of pages. The sonata on-ramp is designed to reward those who are more conservative and may want to entertain various hypotheses about audience segmentation without going deeper into specialized software, cookie dependencies, or other aspects of full-on personalization.

For KPIs, it is intended to increase segmentation, lift conversion rate, and improve confidence in the quality and reliability of segment analytics.

The Symphony

The premise is to begin playing familiar compositions for your users, thereby authoring your organization’s own songbook for what, and for whom, it wants to drive attention to in its digital footprint. Here is where previously conducted user or heuristic research can shine, as it works best with inputs from a jobs-to-be-done (JTBD) or user-needs-focused cookbook of content

recommendations to be offered up in a given component of a certain screen or page of a key digital touchpoint. The idea is to provide a short list of links on a recurring, always-on basis to establish a baseline expectation for users of increasing interactivity and automation. It can be paired with Colin Eagan's zoning approach (Eagan, 2019) to get effect, which will be discussed in detail later in this chapter.

It does not absolutely require recommendations functionality or a recommender software, although that would be a sound basis when scaling it. There are many open-source frameworks for recommenders, and they are generally built-in features to any personalization engine. They are also becoming available as add-ons in marketing automation platforms and some experience management systems.

For KPIs, the symphony is primarily a vehicle for driving engagement and customer satisfaction. It should also be considered a driver of curated traffic (in order to cross-promote more effectively, for example, which is a desirable by-product of this approach as it matures).

Some organizations find the on-ramps compelling as a way of focusing themselves on a clear return on investment (ROI). The best part about them is they don't require expensive software investment and exotic acts of preexisting technology integration. They can be undertaken by small teams or even teams of one: an orchestrated experience with no need of an orchestra.

Taking a more culinary metaphor, by getting a handle on your perso-grammar, operators can focus on the ingredients in the snacks and staples they start whipping up, whether in the "dream kitchen" of a personalization engine software, the cramped galley operation of one's existing warts-and-all tech stack, or a more fleet, food-truck-like operation of a composable digital experience platform (DXP).

Design Is the Critical Gap in a Flawed Personalization Strategy

In the same way in which few content strategies may survive without dedicated ContentOps resourcing, there is no personalization strategy worth its salt without input from a design viewpoint. We're all orchestrators now, and unless you're one of the fortunate few, you may be closer to jug band than symphony.

What's more, a design viewpoint helps broaden the conversation to marketing technology and automation while underscoring the sometimes brute-force content work that can be entailed in more complex orchestrations. These on-ramps deliver interactions that are automated (whether through AI or business rules) as well as customized. Customization is an important flip side to the coin of personalization and a remedy for dark patterns of ethically questionable practices. Oftentimes customization is the best way to offer a more tailored experience to a user, by offering it up as a choice. Regressive personalization is where designers don't offer a way out or an explanation for the tailored interaction.

It's worth mentioning that all personalization "fails" are failures of designing effectively with data. Regressive personalization needs a progressive rethink, a way to tiptoe into more potent, dynamic interactions.

The personalization technology software alone, while wondrous and often exceeding expectations, is not enough. What personalization needs to become a regularly successful practice, more than anything, is clear, rigorous, thoughtful design. At some level, for most organizations this will mean tending to experience debt. This level of commitment may require revisiting a customer journey map. Judging by Colin Eagan's breakthrough work, UX professionals with content subject matter expertise are especially well positioned to make outsize contributions to what design can bring to the delivery of more personalization success stories.

Designing a Clearer Commons

The zone-targeting model gives design and information architecture a proper seat at the table in the conversation around bettering personalization through design. The idea is simple enough, and its plainness is what makes it so self-evidently good and durable: take your personalization program and dimensionalize it across the key real estate that are your customer touchpoints (websites, apps, interfaces, channels, and so on). Eagan's innovation (2019) is to look at personalization like a zoning board looks at building out a civic space. Implicitly, this is sometimes contested corporate territory—or a commons that must be governed and managed.

The joy of the zone-targeting approach, which calls for a rational mapping that makes the business logic or algorithms explainable (and hence reproducible and measurable), is that it can start very small. It can be a matter of isolating a single design element on a single page template or examining how a custom field in an email is used.

Testing this approach is a matter of simply declaring objectives and a clear (and hopefully consistent) grammar for how to utilize the space with appropriate messages for given audiences. Eagan built a telescope that is a microscope, too, and it works at every resolution if you are executing systematically with your personalization program.

The best first step is to get literal and structural and to begin divvying up key components of a customer experience. These primitives form the basis for a design element vernacular that will be automated, customized, or personalized.

Eagan's zoning model takes the 4D hairiness of personalization and reduces it to a 2D concept of real estate to be operated and maintained according to its routine, governed use. There's a lot of conceptual value in that kind of broad classification, and a real estate metaphor speaks especially to the cross-functional challenges of governing corporate digital spaces. Bad UX too often



Jac Rayner
@GirlFromBlupo

Dear Amazon, I bought a toilet seat because I needed one. Necessity, not desire. I do not collect them. I am not a toilet seat addict. No matter how temptingly you email me, I'm not going to think, oh go on then, just one more toilet seat, I'll treat myself.

3:22 AM · Apr 6, 2018 · Echofon

62.1K Retweets 2,400 Quote Tweets 358.1K Likes

Figure 6.2: *Perso-fail* example from a Twitter user complaining about Amazon’s multiple offers for toilet seats. Credit: @GirlFromBlupo.

descends into particulars prior to understanding broader applications and designations of our digital interfaces.

See, for example, Eagan’s model (2019) of bucketing for “tasks at hand” versus “big picture” messaging in personalized settings. Whatever its uses, this method is also a canny way of smuggling in governance by principle and maybe instituting conceptual headroom for organizations to be more measured and balanced in the type of messaging that they deliver as part of a personalized experience.

The Problem with Precision (and Precisionists) in Personalization Programs

Precision in, say, content delivery sounds like a sacred cow of sorts. How could it ever be wrong to give a user the right content?

How about every time a personalization project fails?¹ Every *perso-fail* is instructive in as much as it amounts to indictments of poorly designing with data. Figure 6.2 shows a popular example from Twitter.

This toilet seat is a household durable, not a household perishable, where a retargeting tactic might make somewhat more sense as a marketing “touch.” The perso-fail happens well upstream of the Amazon marketing touch: it’s really an oversight of the rule mechanics for how one system queries another,

¹ A collection is available at the author’s Bucket Brigade (<https://bucket.studio/>), a community site for those who design connected experiences.

probably a marketing automation platform (MAP) and a product information management (PIM) system.

When personalization fails, it is almost always a consequence of overfitting a result, or set of results, that do not apply to the situation at hand. The design remedy is to bias for broader recall of a set of recommendations or possible actions, not a narrower subset, unless there is a strong signal from a user to confirm their intent.

Yet personalization programs, even “successful” ones, often run afoul of vanity metrics in this way, firing off waves of new optimizations, tests, and interventions while not doing enough to attend to broader use cases and needs for a more intelligent, connected interaction. This happens often because of a test-and-learn capability that is maturing and seems well equipped to helm a personalization effort at large as well. Many organizations that seek to scale a personalized customer experience instead drift into the particular, thin-slicing dozens of narrow experiments that cater to edge-case segmentations, ad-hoc promotions, UX hunches, or continuous improvement and optimization of set pages and interactions. These are fundamentally easier and more confident activities, drilling into depths of interaction choice details as opposed to the breadth of interaction choices to get at relevance for the test, the user, and the program owner’s organization.

Don’t get me wrong. Precisionism has its winning qualities: it’s going to be highly attributable, highly defensible, and probably low risk by virtue of affecting so few users. But it is exactly that risk aversion that plays into corporate incentives to play it safe and cordon learnings off to carefully plotted corners of an experience. Too often, precisionism offers an illusion of progress: we’ve run hundreds of experiments!

I’ve seen this pattern play out more dramatically in organizations that outsource their testing and optimization functions and incentivize agencies on a volume throughput model rather than one to deliver against program ROI KPIs.

The other reason for this drift into particulars is that designing for the broad and common use cases is challenging. It is especially grueling for the declaration and prioritization that it asks of business organizations: tell us what you want, in what order. Personalization’s grammar (roughly: if this, then that, for them) is a decoder ring of sorts where large organizations are more comfortable playing in the grey area of hedged bets and the status quo.

I have shared why some of the most crucial project work for my clients is in personalizing whole classes of digital elements, such as web forms or house ad placements or recirculation units: to get at a ground truth for how these function, whether they are dilutive or additive to business strategy, and how to make them more relevant. Information architects like to talk about the difference between breadth and depth, and I’m an inveterate advocate for the idea that depth is too often the business theater of those unwilling or unable to draft out the breadth of their CX—the jobs to be done.

The great *win* of this work is to get a comparative baseline for what that class of element does, and it is the first step toward simplifying and reducing the

clutter and dilution that is taxing your users. It also breeds consistency of practice and measurement. We're not evaluating a call-to-action tout based on how well it satisfies the KPI of a given test: we are able to see how this call to action performs amid all the calls in a given channel or site.

That said, I *am* an advocate for incrementalism—when it's embedded in a roadmap that is structurally sound and truly building toward mastery of the design real estate at hand in your digital and offline touchpoints. Progressive personalization is a framework not only for balancing user value and business benefit but also for generating a consistency of approach and parameters for how an organization personalizes interactions. When you get the fundamentals right, even the particulars flow more smoothly.

Designing for Connected Experience Is a Mission

Simplification isn't simplistic, but it is the end state of any truly successful connected experience. "Simple" requires making the complex clear, and with encouraging results, even subtracting from the content presented. Subtraction is evidence of great focus and one of the great, if rare, feats of superior design, as industrial design hero Dieter Rams has said repeatedly in his career.

As confidence is achieved, intelligent experiences can trim, or visually deemphasize, say, marketing messaging that audience data indicates as unnecessary context or peripheral information to most. It requires someone relentlessly focused on a project of continuous improvement, and it's your mission—if you choose to accept it!—to get to the essence at the intersection of user and business needs. It's a consummate information architecture assignment, one that particularly favors a content specialist.

Connected Experience Demands an Ops State of Mind

Many have worked both long and widely enough in digital to see personalization brave its journey from promise to pratfall to myth and back again. Lately, it's made the run from white whale to white-hot all over again. Personalization is reemerging today with newfound relevance, sophistication, and ubiquity, thanks to a toolbox that ContentOps professionals are uniquely well qualified to deliver or own in organizations.

The real challenge with personalization is executional. It requires systems-whisperers.

The best systems start simple. A winning approach to designing for personalized experience in the warts-and-all reality of your tech stack is about incrementalism. Progressive steps to greater complexity, building on a firm foundation of learnings, will make your tactics more confident and the results more reliable.

The Battle for Breadth in Today's CX Settings

Prioritization is really one of the great issues dogging otherwise good personalization programs that I know and encounter in the field, or I wouldn't stress this principle of getting breadth right, first, before depth. But it's bigger than that. Your assessment of value in personalization may be mistaken if precisionism is driving your decision making. A ContentOps professional understands that componentizing a content offering is key, while it does not need to be totalizing—not all content that could be personalized will or even should be. Eagan's (2019) zoning approach guarantees an atomic-level consideration for content items that must be componentized in order to be part of a progressive path to greater levels of automation or personalization.

For now, I implore you to keep in mind that aside from the formal challenge it may seem to pose, there's real beauty in getting breadth right.

It's not that UX professionals are better communicators or thinkers, though those experienced in working with large corpuses of content or data—think here of information architects especially, though not exclusively—are also frequently adept at what Eli Pariser, author of *The Filter Bubble*, calls thinking in "wide categories" (2011). Again, the ContentOps aspect here is crucial, as breadth is the key challenge of componentizing content altogether. What is the right and most meaningful level of depth and detail?

The ability to land at meaningful generalizations is a recurring motif of effective personalization practice. Across the grammar of any given personalization intervention, generalizability is the ability to deliver repeat value in the content, the audience targeted, the recognizability and usability (therefore, reliability or implied trustworthiness) of the interaction pattern or design elements used, and even when or why the personalized interaction occurs. It's hard to fathom a more useful skill than the ability to think broadly yet accurately about the value users seek from their interactions with your digital products and services.

Practicality in Personalization: Moving to a Performance Mindset

One of the happiest side effects of a personalization program is the focus it puts on content's role in contributing to planned outcomes. Performance is about more than a driver of top-line revenue. It's about better ascertaining the performance of discrete elements in your content offering.

As a certain en vogue perspective in product development has it, users "hire" a website, say, to do a job for them: to solve a problem or answer a question. There are two ways to win at this "job to be done," as design and product wonks know:

- You can supply a satisfying answer, with easy access to ample supporting content or data that will bolster this answer's credibility and additional topical depth for further exploration.

- You can speed the user to that satisfying resolution—or even let them know with similar brevity that you cannot provide the answer in question.

Speed and relevance: practically, those are the performance criteria by which users judge an app session or site visit today. Personalization can help on each front and improve your chances of the site's content, and overall experience, acquitting itself well to the user, therefore reducing effortful interactions and cognitive overload for them but also possibly directly addressing their *ask* in a cohesive, snappy manner. At their best, personalization programs are focused efforts at simplifying the customer experience and speeding their journey, accelerating and scaling the effects of happy paths while accommodating and understanding the need for other paths, too.

Score Performance on a Scale of User Value Creation

For a variety of reasons, it may be beneficial to characterize the value created for a customer. How does your new feature, functionality, experience, or product provide value in their eyes? A balanced personalization program should deliver on a bevy of lower-order needs even if it does nothing life-changing: saving users' time and sparing them cognitive overload already lets them attend to goals they have elsewhere in their daily lives. Bain and *Harvard Business Review* have been advancing the ideas of a hierarchical model of user value (Almquist et al., 2016) and studying its sector-specific variations. You may have user research or testing at hand to validate.

Personalization: Our Most Promising Solution to Information Overload

It's hard to imagine a higher calling for digital professionals than that, but here's one, in all sincerity: Repair the social fabric!

Luckily, it's also the path to an ethically alive and alert regime for any decision making that involves AI or machine learning in one's personalization regime. Explainable AI, or XAI, is fast becoming a legal and ethical operating requirement—just as being able to decompose your automation query logic (more about that grammar below) is essential to monitoring, auditing, and testing those interactions for their verifiable fairness and transparency. Fantastic work in this regard has been underway for years in the recommender system community, such as at the PIReT Ship team at Boise State University (PIReT, 2021), and fairly resembles the ethical table stakes for anyone designing and delivering personalized activities. It all comes down to the grammar that makes any personalized interaction identifiable, repeatable, and scalable. We make those foundational underpinnings strong by leveraging the structural elements of our content.

The so-called attention economy view is that, increasingly, only personalized or algorithmically driven messages will reach us—therefore favoring those who begin orchestrating their marketing and communications in this way. In the early days of the web, this preoccupation was findability; now it's a matter of discoverability and of finding users where they spend their time between, say, their messaging, inboxes, and social media. Amid the demand for discoverability of your organization's products and information, there is a public service of simply, as we've said, saving users precious time. Behavioral scientists talk about IA as a choice architecture, a setting in which nudges can guide users toward the appropriate information "scent" that will drive better outcomes, ideally, for them and the organization.

At its noblest, designing for connected experience is about alleviating and remedying the enfeebling paradox of choice and information overload that so many of us encounter around a multitude of daily decisions, the infinite spread of consumption possibilities our hyperconnected world has laid chaotically at our feet.

References

Accenture Interactive. (2018). *Making it personal*. Retrieved June 20, 2022, from https://www.accenture.com/t20171031T052156Z__w__/us-en/_acnmedia/PDF-34/Accenture-Pulse-Check-Dive-Key-Findings-Personalized-Experiences.pdf.

Almqvist, E., Senior, J., & Bloch, N. (2016, September). The elements of value. *Harvard Business Review*. Retrieved June 20, 2022, from <https://hbr.org/2016/09/the-elements-of-value>.

Eagan, C. (2019, January 17). UX in the age of personalization. *A List Apart*. Retrieved June 20, 2022, from <https://alistapart.com/article/emerging-ux-role-in-personalization/>.

Gartner. (2019, December 2). *Gartner predicts 80% of marketers will abandon personalization efforts by 2025*. Retrieved June 20, 2022, from <https://www.gartner.com/en/newsroom/press-releases/2019-12-02-gartner-predicts-80--of-marketers-will-abandon-person>.

Pariser, E. (2011). *The filter bubble*. New York: Penguin.

PIReT. (2021). *People and information research team*. Retrieved June 20, 2022, from <https://piret.info/>.

Soegaard, M. (2021). *Hick's law: Making the choice easier for users*. Interaction Design Foundation. Retrieved June 20, 2022, from <https://www.interaction-design.org/literature/article/hick-s-law-making-the-choice-easier-for-users>.

CHAPTER 7

Localization and Content Operations

Loy Searle

Localization Overview

The Localization Profession

Localization is a global industry that focuses on translating content and other types of deliverables from one language to another with an eye on making the language output either globally appropriate or locally appropriate. Sometimes, localization and translation can be used interchangeably. In this chapter, I provide an overview of localization (and its related terms and practices), and then I focus on the interesting intersections between Content Operations and localization.

The most obvious intersection of localization and ContentOps is their focus on the written word—its quality and usage—as well as language and their integrated processes, handoffs, and shared publishing systems. The localization business is largely outsourced. While there are full-time equivalents (FTEs) employed throughout the process, translation work is frequently performed by freelancers. Consequently, a big focus of the industry is budget availability, optimized processes, and scalable technology.

Before we go too deep, a good place to begin is with some key industry definitions. You will see acronyms related to functions that are commonly used within the profession. Each acronym is determined by the number of letters between the first and last letter of the word. For example, *translation* has 9

How to cite this book chapter:

Searle, L. 2024. Localization and Content Operations. In: Evia, C. (ed.) *Content Operations from Start to Scale: Perspectives from Industry Experts*. Pp. 109–124. Blacksburg: Virginia Tech Publishing. DOI: https://doi.org/10.21061/content_operations_evia_9. License: CC BY

letters after the T and before the final N. (A fun party game is matching people to their acronym. For example, my name, Loy, is L1Y according to this abbreviation system.) For a history lesson, you can Google the origin of i18n. If you find yourself writing *internationalization* many times a day, you will quickly understand why these abbreviations have become so deeply ingrained within the industry. Figure 7.1 includes some definitions of fundamental terminology for the localization industry.

Translation (T9n) is taking a concept or an idea in one language and expressing it in another language. A proper translation requires that you really understand what you're expressing conceptually, including its context, audience, and purpose. A bad translation is a literal word-for-word expression from the source language.

A translator's work starts with inherited source content. The translator must make sense of someone else's words to do their job. Consequently, they will find gaps and problems in the source content, so a good feedback mechanism between content and localization is essential.

Localization (l10n) is often considered synonymous with translation, but it involves more than that. (Actually, translation is generally considered to be an aspect of localization.) Localization is perceived as the broader category. To localize implies that you are going a step beyond translation to make your deliverable appropriate for a specific locale (e.g., country, state, region), industry, or work. The process of localizing can be intense!

Also, to add to the fun, localization is what the *overall industry* is called. You can see why translation and localization get confused.

While we usually think about localization as making something appropriate for a specific language and locale combination, often the work supports many locales that use a language. For example, Spanish is spoken in many countries, but the local terminology (and in some cases, even grammar) can vary by country. A common practice in the industry is to attempt to standardize the Spanish in a way that it can be understood in many locales. Latin American Spanish or Neutral Spanish are two approaches that allow an organization to address many locales with one language.

There are two focus areas for localization. One is language localization—which is often a synonym for translation. The other area is product feature localization. In some industry sectors, there is no distinct difference between these two areas. If a product works the same globally as it does locally, there may be no unique feature localization requirement. An example might be Gmail—which does work the same way in most countries.

However, in highly regulated industries (such as finance, banking, and accounting), these two types of localization are often differentiated. The reason for this differentiation is that a translator does not have the ability to make a product work properly locally—they can only change the words or labeling. If a product must behave differently, perform different calculations, or work

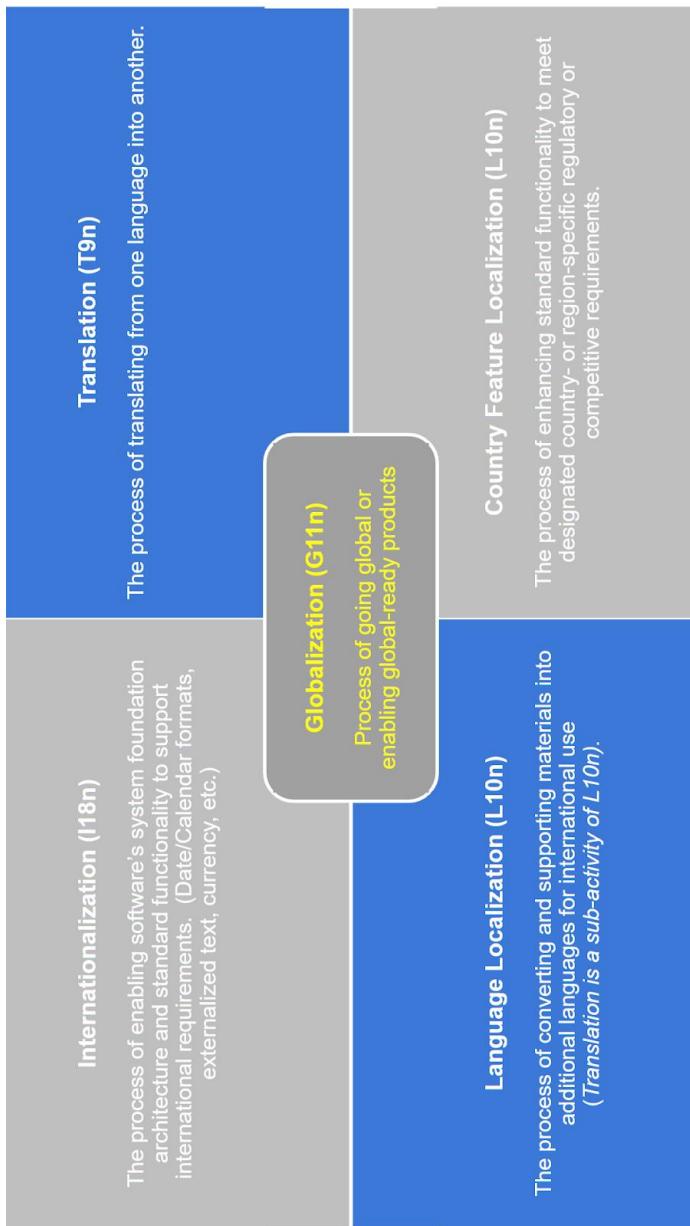


Figure 7.1: Common localization industry definitions. Credit: Loy Searle.

differently in another market (outside of what is typically handled through internationalization), the product will need to be designed and coded for that different local functionality. For example, tax rules can vary by country, state, county, and city. There can be multiple types of taxes. The work to adapt a financial product requires product feature localization that goes beyond label translations.

Internationalization (i18n) is the practice of ensuring that your products will work globally. There is a long list of items that fall within this grouping, but here is a standard summary: date formats, calendar formats, externalized text to enable translations, country and city codes, labels and address formats, and currency and number formats, among others.

Globalization (G11n) has a couple of different meanings. It can mean the geopolitical global redistribution of work, such as in manufacturing—where we've seen an industry shift work from the home country to another country as a result of cheaper labor or resource costs. This is not the definition of globalization that I will be covering.

Globalization, as it relates to the localization industry, is generally considered to be the broadest term used to cover what we do within the industry. When a company is planning to take its products global, there are many tasks that must be done to make this move effective. The practice of completing all the tasks involved in making a product work globally is often called *globalization*. A product can also be “globalized,” which means that all these activities have been done. The globalization tasks that I will address in this chapter will be limited to the ones that typically impact the localization profession.

Globalization frequently includes translation, localization (language and product features), and internationalization. Within a large company, there may be teams that specialize in each area, and it's common for them to be organized together or, at a minimum, to work very closely together daily.

So, localization and globalization teams can be the same—often because organization names can have a long shelf life, while team members can fluctuate over time. In general, a globalization team almost always includes ownership of most of all these components. Regardless of a team's name, there will likely be people in roles responsible for accomplishing goals in each of these domain areas.

How the Client and Supplier Sides of the Industry Work Together

As I mentioned above, localization is largely an outsourced profession. Translators live all around the world—most often in their native country. They are usually freelance workers who are employed by translation vendors or suppliers, or, on occasion, directly by customers. The business of localization requires many roles for support. Some roles are similar on both the customer and supplier sides of the business.

The supplier side of the industry is broken down by very large multi-language vendors (MLVs), single-language vendors (SLVs), and freelancers. MLVs typically engage with many SLVs or smaller MLVs as well as freelancers. A large part of this side of the business is finding and managing those resources. The real work of “translating” is often done by trained linguists who likely specialize in certain types of translation work or in a specific industry (medical, software, or marketing, for example) and may work for one or more vendors supporting multiple clients.

On the vendor side, a project manager (PM) handles customer submissions through the translation, review, desktop publishing (DTP), and testing processes. On the client side, a PM works with requests coming through from internal stakeholders, through to the vendor—including pricing, timing, and so forth—and back through review, approval, perhaps testing, and delivery to the internal stakeholders. As a result, these roles have some similarity but can be experienced quite differently. Both roles require exceptional organization and communication skills, attention to detail, and technical understanding of the deliverable types.

Localization is very similar to a supply-chain business. Imagine if Amazon delivered awesome translations instead of sending packages. There is a similar level of complexity with work flowing around the world on tight deadlines: aligning people who can do each part of the complex process and managing estimates, payments, different currencies, different global schedules, special demand over holidays, global business, occasional misdeliveries or returns, quality challenges, and customer feedback. Because of the complex nature of the industry, automation, process, and scalable systems are critical to daily operations—as are amazing people who understand the work. It’s a fun profession that has a tightly tuned sense of time everywhere.

To normalize localization work across clients, vendors, and freelancers, the unit of work in this industry is frequently the word. Cost is often a blend of word rates and hourly rates for tasks that are less measurable. It’s important to not just consider word rates as the core cost, as every project typically involves PM work as well as perhaps review, testing, and other responsibilities. That said, this is a highly budget-conscious profession. Everything is counted, measured, estimated, and validated. So, while it’s a bit like Amazon from a supply-chain-management standpoint, it is also like a procurement function, where every dollar must be well managed. This is true on both the supplier and client sides of the business.

How Localization Works

Localization is a blend of technology, process, and internal and vendor-managed people. It’s a profession that relies heavily on process and technology. To best understand how localization works, I’ll dive into the technology

landscape that intersects most with ContentOps, cover the standard processes, and then review the intersection points between localization and ContentOps.

Technology and Translation Memories

Localization leans heavily on third-party technology solutions. In the past, these tools were often highly customized on premises. Today, most providers have shifted to cloud-based applications. Following are the technology elements from the localization industry that intersect with Content Operations:

- Translation memories (TMs). The core capability within a TMS or translation workbench that powers translation.
- Translation management systems (TMS). The translation workflow system.
- Terminology management (term bases, or TBs).
- Query management and issue/bug reporting.
- Machine translation (MT).
- Publishing tools.

Translation Memories

The one capability that is unique and fundamental to localization is translation memories (TMs), a database of all the sentences or text strings that have been translated in the past in both the source and target languages.

When a project is analyzed the first time it gets translated, it is added to the translation memories. Both the source and the target translations are saved and matched. The next time the same project comes through, maybe with a percentage of changed material, that project will be compared (analyzed) against the existing memory, and only the changed segments will require translation. All segments will appear and can be edited, but the translator can quickly work through just the new and changed segments. Translation memories minimize rework and reduce costs by ensuring that time and money are not spent on unchanged or only lightly changed segments.

A new project will come through with no or a very low percentage of matches. If it returns with updates from a new release, there will be a percentage of “fuzzy matches”—indicating that strings (or sentences) partially match the existing memories. If part of a phrase is changed, it will become a fuzzy match. A project with many fuzzy matches will be cheaper to translate than a new project. If a sentence is exactly the same as a previous translation, and nothing changed when it was resubmitted to translation, this sentence will be a 100% match. If the context is also exactly the same, and the sentence has not changed (as in the same file

resubmitted in the next release), it will be considered an ICE (in-context exact) match. ICE matches are typically free (100% matches are very inexpensive). The lower the fuzzy match percentage, the higher the word rate for the translation. Therefore, savings are incurred through always analyzing every new and changed file against the TMs to “leverage” from the previous translation.

A new sentence can benefit from leverage if rigorous standards are used in the source files. These standards can include repetition of commonly used phrases, as a new file can come in with some fuzzy matches.

Changing source files once they've been translated can be problematic. Gratuitous change should be avoided if possible, as even changing punctuation standards will change segments. My best guidance is this: If something needs to be changed, change it. A penny saved in translation living with a poor source language sentence isn't worth it. Make your content better, and those improvements will flow through the language translations.

From a process standpoint, an ICE match will likely not be reviewed in the translation process, as it didn't change. Fuzzy matches typically are reviewed within the translation process. If the original file was not properly reviewed the first time, it's possible that a poor translation within the memories can be reused. A good common practice is to spot-check even ICE matches to ensure that the overall quality of the deliverable remains high.

Translation memories are an investment, as ensuring that quality is high will improve the quality of future deliverables. Do it right the first time, and less money will be spent downstream.

Translation Management System (TMS)

A translation management system is the first and most important tool that will be purchased by your localization translation team. It is a workflow solution as well as a translation workbench that supports translation activities such as translation memory management, terminology management, and sometimes query and quality management. It may also have basic financial handling capability (preparing quotes, estimating, and handling invoices, for example). Lastly, it will include some reporting capabilities.

Today, most TMSs are cloud-based software offered as a service solution, but many large companies are still using on-premises solutions that they likely customized years ago. A TMS is a highly specialized system that is designed to support client-side work initiation, the transfer to vendor-side for translation, review, and other post-processing work, then back again to the client side for linguistic quality analysis, testing, subject matter expert reviews or sign-offs, and then delivery. A robust TMS can handle different types of work—from user interface to marketing content, to help manuals or training materials, to video content.

Terminology Management

Another unique capability within translation is terminology management. Almost all translation workbenches have integrated terminology management, and there are third-party terminology management systems. Terminology is important because translators are generalists and may not understand a specific company's or industry's terminology. This enables a translator to do an acceptable job, even if they are not deeply experienced with a given company's content.

If terminology is not defined, a translator will either research it (which is time-consuming), guess (which is a risky bad practice), or request a definition while they are translating the work (which may or may not arrive in time). Therefore, term definition is ideally a proactive task within translation, but it is often also reactive.

Writer glossaries can be imported into terminology databases (or term bases) to provide a good start for translators. As translators do their work, they will identify terms from the term base to improve consistency, or they can define or flag them for future definition. Term bases include both the source and the target terms. They can be simple equivalents or much richer, depending upon the investment and need. Key terms will contain more information than a simple writer glossary will, such as a “concept” or what it means, perhaps an example, gender, associated domains, and more. Because terminology consistency is part of the translation process supported by automation, translation terminology handling is often more consistent than the source content, which may lack these checks. Also, terminology can be aligned across many types of deliverables for even greater consistency.

Both translation memories and term bases are core elements of the translation process. They enable translation organizations to truly align voice, tone, style, product, and feature names as well as terminology across a company. Obsolete names are also tracked to ensure that they are not used. This is difficult to achieve in a source language when many teams tend to create content.

Industry Publishing Tools

Localization typically does not own its publishing tools. Localized content is typically published into whatever tool the request comes from (GitHub, content management systems, knowledge management systems, learning management systems, community sites, and so on). This means that localization teams need familiarity with a wide range of tools to validate deliverables, perform in-context review, and automate handoffs.

Machine Translation (MT) or Neural MT (NMT)

Machine translation is just that: a machine taking a stab at translating from one language to another, based upon grammar rules and common terminology. Standard MT is trained via rules, and NMT leverages machine learning.

Machine translation has become ubiquitous in the translation industry. The quality levels have improved, and with NMT—which is MT that learns—the naturalness of speech has come a long way from the early days of BabelFish. MT is a common capability within most translation teams today. It is used for “gisting” as well as preloading a translation with an MT suggestion. MT is rarely used for customer consumption unless it is “post-edited” by a human reviewer or unless a user directly requests it. Because it is machine-generated, tricky or highly technical terminology, colloquialisms, and poor-quality source material can all have a negative impact on the target language quality. For this reason, the best practice is to train the MT so it improves through feedback via post-editing and a high volume of good examples.

MT quality can vary by language and by the nature of the content being translated. If you have used Google Translate, you have used an NMT solution. NMT is trained by processing huge volumes of language strings. Longer-tail languages tend to perform more poorly, as there is less data to train the engines. Linguistic distance from the source language is also a factor that impacts quality—the greater the distance, the lower the starting quality. Western languages tend to be closer, so the results can be surprisingly good. MT quality is measured by the distance between the machine translation and the human post-edit. Few edits imply good quality.

When and how to use MT requires careful thought about quality, the audience, the languages, the shelf life of the content, and the urgency of an immediate translation. It’s a highly useful capability, but across the board, it does not yet have an edge over human translation.

Raw MT

MT is considered raw when it is served up without a post-edit or human review. MT is rarely served up raw unless one or more of the following is true:

- A user specifically asks for it (you request Google to translate a screen in another language). Your quality expectations should be appropriate.
- The MT has been well trained, the quality levels have been vetted, and the deliverable is appropriate for MT.
- The deliverable has a very short shelf life (e.g., a blog or newscast transcription). Timeliness matters more than quality.

- The deliverable's quality is not important (either to the user or perhaps to the company that hosts it).
- Gisting is sufficient (you need just enough information to generally understand something; perfection is not required and would be a waste of time and money).
- Providing a pre-MT suggestion is part of a translation process to save time.

If you do use raw MT, you should tell your readers and ideally let them make the decision to consume it. This will spare you a lot of quality problems (and some embarrassment).

Query Management or Issue/Bug Tracking Tool

Translators often work with content out of context, and they often translate material that is new to them. Consequently, they will have questions when translating. A query management system helps handle these queries and typically integrates with a bug tracking system and terminology management. If forty different language translators have the same question when translating a project, imagine how happy your subject matter experts will be to triage forty versions of the same question!

Query management emerged to help streamline questions that would be applicable to many languages or that might reveal a bug or a need for new terminology. Queries could also be language-specific. Answers to queries need to be quick to catch the project while it is still in flight. Queries are first reviewed by the vendor; then they go to a PM or a linguist on the client side—depending upon the type of query. Next, they might go to a subject matter expert (like a writer) for their help. Lastly, they may go to an engineer if a bug is determined to be the cause of the query. Some TMSs have query management capabilities, and bug systems can be customized for this purpose as well. There are also purpose-built tools for this work that can integrate with a TMS, terminology system, and issue management system.

Translation Projects and Process

The common unit of work within this field is a “project,” which can consist of many types of files or deliverables that are parsed by number of words. The type of file or work being done can vary wildly, and here are some examples of common types of work that is translated:

- Marketing sites and content.
- Product user interface (UI).

- Physical assets, such as brochures, books, banners, stands, and products.
- Videos, movies, and games.
- Help and user content, support, community, instructional content and tutorials, and presentations.
- Legal documents.
- Transactional content or data.

Localization teams typically handle broad types of deliverables. Some deliverables will be predictable daily activities; many will be ad hoc. Thus, localization teams must be comfortable working with all content types and technology—and their vendors must be equally proficient. Each deliverable type will have a defined process that can differ by audience, quality expectations, time available, and budget.

Almost every one of the deliverables above will undergo more than just translation within the localization process. There will usually be reviews and perhaps testing, in-context validation, DTP work, and so on. Video work or custom work can follow a different process, and it's not uncommon for a translation team to have a dozen or more commonly used workflows. Figure 7.2 presents a standard translation process coming from a client through to a vendor.

This process assumes no hiccups, rejections, late deliveries, or mishaps. When reviewing the process, it's important to consider the nature of the business requesting the work. Some businesses have thoughtful, predictable deadlines for much of their work, while others function in hyperdrive. Some work requires impeccable quality with multiple checks. Within a business, not all content is equal: some work may be highly visible, important, and durable, but other content may have a very short shelf life or require fewer quality checks. Indeed, while fifty companies may have a very similar high-level, the experience can be wildly different depending upon the nature of their business, pace, and deliverables.

Localization teams are wired to build scalable processes and technology. This is not just nice to have—it's survival! Imagine a medium, three-person-sized project being translated into 70 languages. This project might be sent to 210 translators all around the world, with another 70 reviewers. Their availability for work requires engagement and planning. If work is delayed, resources might be committed elsewhere, resulting in delivery delays. Now imagine that there are 20 of these projects coming through every day from different teams. Ideally, there is a continuous flow of work to the vendor from the customer, so a ready set of resources familiar with the company's requirements will be on hand. If not, recognize that it takes time to build, nurture, and mature a resource team. Ultimately, it's a people business underneath all the automation and process optimization. The quality of the translation depends not just upon a great process and tools but also good source materials, clear guidelines and standards, sufficient time, and most importantly, knowledgeable, committed, and highly capable translators, engineers, and PMs.

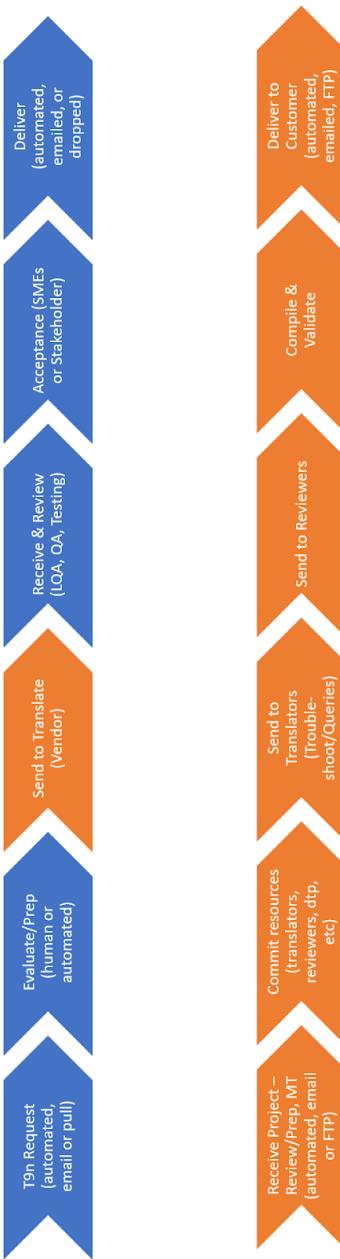


Figure 7.2: Best practice translation process. Credit: Loy Searle.

Translators often work for many customers through one or more vendors. They are usually paid by the word—which means the faster they work, the more they get paid. One of the best ways to ensure a quality translation is to verify that the source is clear, clean, concise, and unproblematic to process. Imagine 280 people in different countries trying to make sense out of a poorly written English text, and then imagine what the language quality will look like in their end deliverable. “Garbage in, garbage out” was a phrase tailor-made for this industry. You cannot ensure a great translation with a great source, but you can ensure a poor translation with a poor source.

Translators are also generalists. If you have a vendor who retains the human resources working with your account, their knowledge of your products, voice, tone, style, and terminology will grow over time—so your translations will improve over time as well. They also don’t have a direct line to the creator of the content that they are translating. Most vendor and client teams have “query management” processes and tooling to streamline fast responses to translators while they work to ensure that they don’t have to guess. Often, the answer misses the translator but catches the reviewer, who can make a change, if needed, based upon the answers given. These queries are likely to first be triaged by the vendor and then by the client-side localization team, which may then reach out to internal stakeholders for answers. It’s a game of telephone via workflow automation—and is always a race against the clock to meet the deadline for when the project is due.

The best results come from well-integrated technology that is optimized to minimize human handoffs of projects on both the vendor and client sides. A strong TMS optimized to automate predictable handoffs reduces human errors and idiosyncrasies and minimizes process time. The best results come from a rock-solid, tight process that is predictable, reliable, understandable, well documented, and (wherever reasonably possible) automated. Your process should provide good information to the translators in the form of glossaries, references, in-context source content, responsive query management, quality assessment, and feedback.

I emphasize that localization is a process-heavy industry, because the goal is typically to take whatever is thrown at us in one language and quickly get that translated into many languages, make certain that it reads beautifully, works, and is tested or validated, evaluated, and often even personalized or customized for a specific market or type of customer before it is delivered as expected, without too much drama or fanfare visible to the requesting stakeholder. Underneath it all are the translators, reviewers, and testers doing great work, PMs on the client and vendor sides keeping it all flowing without a hitch, engineers keeping the tools working and automating as new work types come in, and engagement skills at all levels to work smoothly across different cultures, time zones, and companies.

Does Contact Strategy Impact Localization?

Localization processes can handle almost any content process that is capable of externalizing text into a common format or has application programming interfaces (APIs) compatible with standard translation formats. That said, an authoring process that separates content from its published structure can save costs in translation. If content is reused in multiple types of deliverables and the markup is not changed by each system that reuses that content, those content strings will flow through translation with a high fuzzy matching score—or, if unchanged, a 100% matching score.

Content reuse can create savings in translation—as much as it does in writing. However, it can be taken to levels that are useless or sometimes problematic within translations. For example, partial sentences or snippets strung together or mixed with standard text can flow through translation, but they cannot function the same way within published translations as they do within the source content. Each snippet simply becomes part of the whole translated string and must be divorced from the original snippet in the target language. This is because grammar differs across languages, and partial sentences will not have the same grammar structure in other languages. These simply don't work in other languages. If the CMS supports conversion of snippets to regular text, this won't present a problem within translation.

Content workflow, change management, and version control also impact localization processes. If content is dynamically changing and is also submitted to translations dynamically, version control and workflow steps that include translation handoffs and returns must be in place. Without them, there will be no way to ensure alignment between the writing and translation processes. If these capabilities don't exist on the authoring side, a translation cutoff will need to be determined, and translations will likely not be done dynamically but rather at set dates—catching what's changed since the last iteration, and often not aligning with the source language until the brief period of the next catch-up project.

Another area that is often overlooked within content strategy is alignment of user experience and voice, tone, style, and personas across all—and I do mean all—content types. There may be additional roles with some deliverables, or more focus on a specific demographic for one product area, but there should be an overarching goal to show up in front of your customers sounding like the same self-aware company. This matters because content will be reused by people, and those clashes will show up in the source language and cause rework of duplicative content. Also, your customers don't care about your organizational structure. They consume content across organizational silos and experience those organizational disconnects. Lastly, members of a localization team see every one of these disconnects, and they may be able to normalize some of this in each target language—but they can't prevent the disconnect that typically exists across all the content organizations.

One great practice that can improve the overall content experience is authoring content through the customer process engagement lens, rather than by functional area. If a customer starts in marketing—looking at the product, then purchasing it, then training on how to use it, then starting to use it, having questions that might be answered by help or community, or having issues to resolve with support, they would prefer a more seamless process that is aware of what came before and what comes next and also demonstrates some consistency of voice, tone, style, and user experience. This type of design-driven process typically requires communication and sharing between the different content systems, plus common taxonomy and metadata. This approach should also include localization to ensure that country, language, and local content can be factored in as well.

Do Content Operations Impact Localizations?

Most localization teams aspire to highly automated handoffs, streamlined processes that leverage standards, predictability, and tight communication. When our partner organizations share the same motivation, amazing things are possible. Handoffs can be fully automated, and post-production pre-publishing rework can be minimized or eliminated. Less internal localization labor is required to fix or compensate for poor structure, processes, tooling, publishing, and planning from the requesting content team.

If ContentOps are well managed with tight processes and technical integrations that collaboratively engage and factor in localizations, savings in time, cost, and process friction are a given. Content and localization teams are often closely integrated and perform as happy partners. Both organizations usually love structure, process, and robust enterprise tooling.

One of the biggest challenges that localization teams encounter is when content teams seek to overstep their domain and try to solve localization problems within their content silo. A content team that introduces a separate localization tool or MT solution is not doing the company or their localization team a favor. They are instead introducing friction, doubling the processes of the localization team, and likely increasing costs and delivery times while decreasing quality outputs. In short, don't shop for localization tools. Work with your localization team on how to optimize your team's processes and handoffs to work most efficiently within the localization team's processes.

Bundled authoring solutions that include translation tools are only a great thing when the localization team also wants to use those tools and joins in the process of selecting the right collaborative solutions with a strong partner voice at the table. A last point on this: If your content will be translated, you'll need to include your localization team as an equal partner in the technology selection processes. If you purchase a CMS that only works in English, you won't be solving for your company or your customers. If search only works

in English, or global fonts are not supported, keep shopping. If there are no APIs, or if your tool of choice suffers from lack of interoperability by creating its own translation tools within the product ... run, don't walk, away from that solution! It will be gone within 3+ years, and the folks who chose it will likely be gone too.

Few companies today manage content in a holistic manner. Because localization can't impose order upstream, they focus on API integration of standard industry tools so that translations can all be processed in an efficient manner through their localization processes, avoiding one-offs and manual handoffs. Content organizations are often organized functionally, indexing toward their organization and audience demands, and therefore they will use different types of content management tools:

- Structured content or Darwin Information Typing Architecture (DITA) CMS for technical publications.
- Learning management systems (LMS) for training organizations.
- Knowledge management systems (KMS) for support organizations.
- Marketing CMS, digital asset management systems, website management, and more for marketing organizations.
- File management systems for engineering teams.

Because all these systems typically handoff to translations, the process requires APIs, plus a baseline of global maturity, in each of these content-related systems.

Summary

The localization profession is a rich field that leans heavily on process, technology, and outsourced services managed by internal teams. While both the localization and content professions focus on delivering the right words for the right role and purpose, localization tends to focus more energy on scalable processes than the content field does. That said, there are similarities and dependencies with Content Operations and strategy. Most dependencies and similarities are tied to shared deliverables, joined processes, and technology to transform and deliver source-language content into localized content.

CHAPTER 8

The Technology that Supports Content Operations

Patrick Bosek

Content Operations intrinsically relies on systems and tools. In the purest conceptual sense, ContentOps could be separated from technology, but practically speaking, the two are deeply intertwined. The way in which you implement your technology will be driven by how you implement your processes, and vice versa.

This chapter explores the technology that supports ContentOps and some of the choices that it presents to you. The intention is to leave you with a broad understanding of what's available and how to use it in your ContentOps ecosystem. It is worth noting that in this chapter I will avoid making any specific tooling recommendations.

When starting a ContentOps project, especially before arriving at technology selection and implementation, you'll want to have a strong plan. In this chapter, I'm introducing a battle card¹ as a helpful way to organize the high-level aspects of this plan. Later in the chapter, I'll talk about filling it in.

¹ “Battle cards are internal documents used in the sales field that compile key information and scripts about specific products or services, the market in which they're sold, brand competitors and target customers. Sales professionals can quickly reference battle cards during calls and presentations to locate important statistics, figures, product specifications, clips and other details that may strengthen their pitch.

How to cite this book chapter:

Bosek, P. 2024. The Technology that Supports Content Operations. In: Evia, C. (ed.) *Content Operations from Start to Scale: Perspectives from Industry Experts*. Pp. 125–145. Blacksburg: Virginia Tech Publishing. DOI: https://doi.org/10.21061/content_operations_evia_10. License: CC BY

Table 8.1: Sample battle card for selecting and implementing technology related to ContentOps.

ContentOps Implementation Battle Card			
Channels	•	Content	•
Collaborators / Constituents		Key Business Objectives	
•		•	
Major Challenges		KPIs	
•		•	

Before We Start Building ...

Let's start by asking some key questions and going through some basic planning. This planning work will be the basis for all of the technology decisions that you make downstream. Some of these questions will be readdressed at the end of the chapter, when we look at filling in a project battle card. However, it's good to ask them up front so some of the thinking is fresh in our minds as we consider the technical details of a ContentOps infrastructure.

What Channels?

Here, you should gather a quick list of known near-term and likely future channels for content publication. At this point, it's important to know which channels are **definitives**, but also to have some **maybes** in mind. Without the maybes, you can get too locked into the **definitives** that are right in front of you. One of the recurring themes during the implementation process that I

In addition, battle cards may include tactics to overcome sales barriers, such as customer objections, competitor comparisons and more” (Indeed Editorial Team, 2021).

recommend will be avoiding blinders and short-term thinking—and looking toward future channels is an important part of this effort.

What Content?

Outlining and documenting what content is in scope is an important aspect of planning because it allows you to focus on the solution space and know when that changes. When another group with a different kind of content wants to start leveraging the ContentOps ecosystem, this possible collaboration is a discrete event that can trigger the proper planning and implementation.

Which Collaborators?

The people involved in this implementation can't be an afterthought; indeed, they must be one of the first things you think about. Their objectives, preferences, digital territory, roles, background, and physical location all play significant roles in how they participate in the ContentOps ecosystem. Later in this chapter, I will present some planning strategies that dig into this a bit more.

How Much Consistency Do You Need?

This seems like a fake question, but it's actually something to consider. Different types of content need different levels of consistency. Informal, single-channel content (like many blogs) doesn't require much consistency. It's fine if different authors take different approaches to visuals, layout, voice, and markup. On the other end of the spectrum, instructions for use (IFUs) for medical devices require tight controls around consistency of source content. For more on this, you can read chapter 5 in this collection, where Rahel Bailie and Carlos Evia introduce levels of robustness for ContentOps implementations.

Additionally, it's very common to have multiple content types that require different levels of attention to consistency. Therefore, it's important to catalog this before our next phase.

How Much Structure and Semantics?

When entering into the information architecture and planning phase later on, you'll need to do some concrete work on defining the required structure and semantics. For now, it's helpful to start thinking about the “is-ness” of your content and how important it is. An easy example is *codeblocks*. Is semantically

defining a block of code as a *codeblock* (this content *is* code) important? This one is easy, because almost every organization that has code in its documentation will answer “yes.” Other parts of your content will be trickier.

Well-defined semantics unlock tons of potential for efficiency and user experience improvements, but they also create overhead for those maintaining the content. More semantics aren’t inherently better—it’s a balance.

How Much Metadata Do You Need?

Metadata, as we’re using it here, is *on-object* metadata. This will sometimes be referred to as tagging or classifying content. This is different from content semantics that are *in line* with the content, even though semantics are also technically metadata.

Generally, more content means that you need more metadata, but it’s not always that simple. Another major factor in determining how much metadata you need is how much *similar* content you have. When different pieces of content are very similar but distinct in important ways (e.g., a repair procedure for two different versions of a product), it becomes harder for standard search mechanisms to identify the correct piece of content without the user specifying some additional information to filter results.

The Technical Goals of a ContentOps Stack

- Interoperability—accessible via an Application Programming Interface (API) and friendly
- Predictability
- Governance
- Auditability

The core technical goal of a ContentOps stack is to automate things. This automation can take many forms, but at the end of the day, almost everything that you implement has the goal of automating something. As you’re considering the implementation, one of the best questions to ask is “What are the most important things to automate?” This question can inform which technologies you choose, because different technologies excel in different types of automation.

Content Supply Chain vs. Content Ecosystem

It’s important to make the conceptual distinction between your ecosystem and your supply chains, because confusing them can lead to poor architectural

decisions that result in future limitations. You’re building an ecosystem, and inside it there will be supply chains. Every set of systems and processes that take content from origination to consumption is a supply chain—with the awareness that content is never “done.” It has to be maintained, reviewed, updated, or retired. The totality of the connected technologies is your ecosystem.

The most common case of confusing supply chain with ecosystem is when a group views its entire implementation through the lens of a single target supply chain. In this case, the group will often choose the least expensive technologies to support that single supply chain. When future supply chains need to be added, these technologies often fail to meet the new requirements.

Architecting Confidence to Achieve Agility

The ability to move quickly is intrinsically linked to how confident you are in the outcomes produced by your systems and processes. This is fundamental to ContentOps implementations, but often it’s not top of mind.

Confidence in the content and the user experience of the systems delivering that content is what allows you to publish content with agility and at a rapid pace. Every bit of uncertainty in any element of the final content product introduces quality control steps into the process that increase time to market and total cost to deploy. A good ContentOps implementation provides an organization with a method of producing all necessary content products with the least cost and time between the decision to publish and the content being consumable. To achieve this, confidence needs to be actively considered throughout the planning and implementation of the system. You have to actively architect confidence into your ContentOps ecosystem.

Offered vs. Published

One of the key aspects of a content supply chain is how the source content actually becomes user-facing content. Broadly speaking, there are two different ways this can happen: offering and publishing.

Offered content refers to the process of making content available over an API and allowing apps to retrieve it as needed. *Published*—or delivered—content, on the other hand, is when a system builds content into a static package that is moved to another location independent of the system that built it. Both methods have their purposes, and there are situations where the line between these two methods is blurry.

Benefits of Offered Content

- Makes at-scale personalization possible (for a deep dive on personalization and ContentOps, please see the chapter from Jeffrey MacIntyre in this collection).
- Automates the process of propagating changes.
- Simplifies integration with other systems, enabling omnichannel content workflows.

Benefits of Published Content

- Creates artifacts.
- Works well with standard content delivery networks for maximum performance.
- Can be good for offline or other situations where portability is important.

The Components of a ContentOps Stack

At a very high level, every ContentOps stack has four main components:

- (Structured) content
- Content source
- Content services
- Delivery apps

ContentOps Supply Chain

Each of the components of a ContentOps supply chain (as in Figure 8.1) has a set of functions and responsibilities.

Expanding into each of these major components, we get something that looks more like Figure 8.2.

Next, we can add some surrounding systems to start visualizing an ecosystem, as in Figure 8.3.



Figure 8.1: Conceptual primary components of a ContentOps supply chain.
Credit: Patrick Bösek.

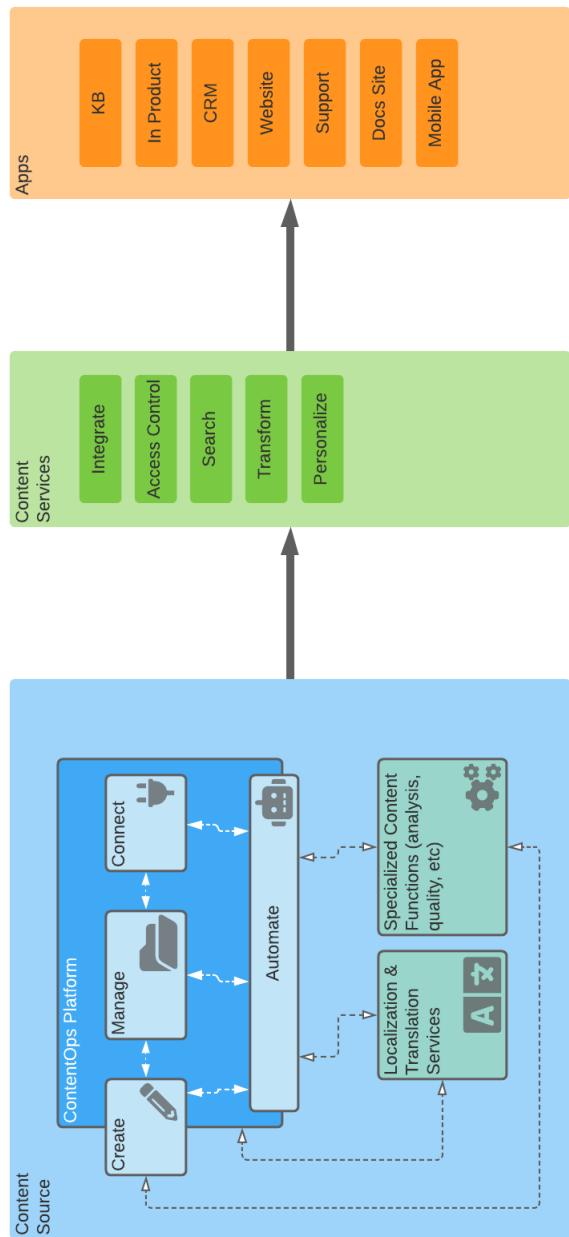


Figure 8.2: Base ContentOps infrastructure by supply chain component. Credit: Patrick Bosek.

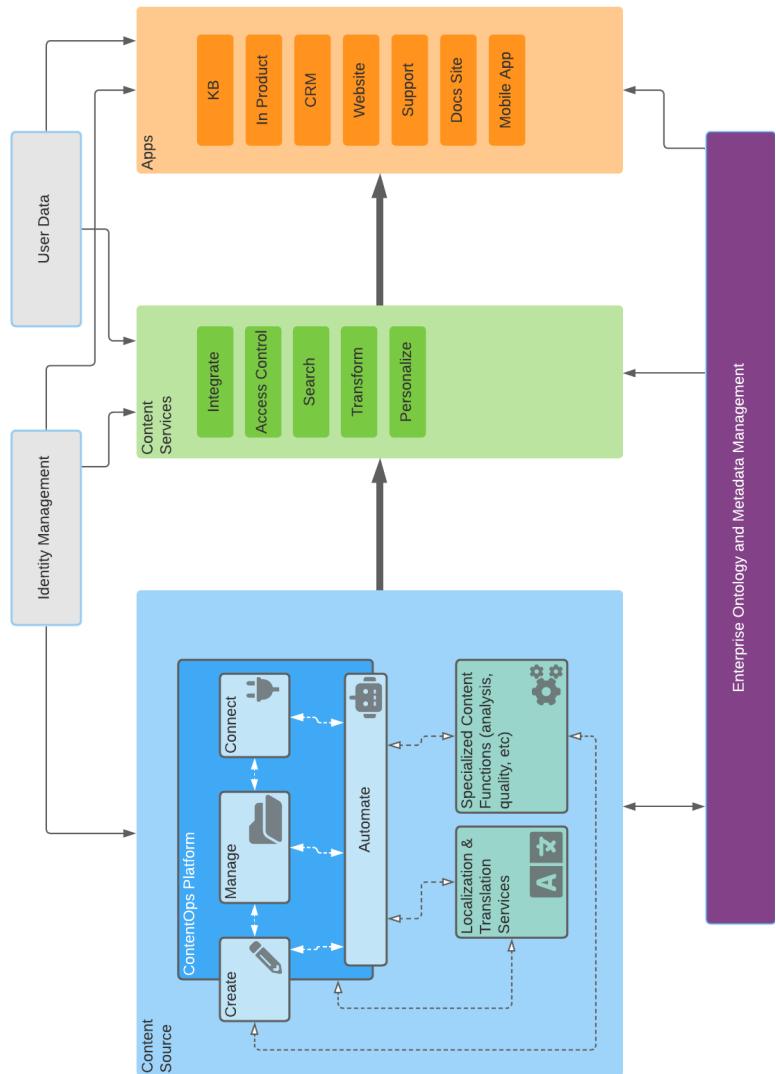


Figure 8.3: ContentOps infrastructure with common connected systems. Credit: Patrick Bösek.

Content Source

The content source is really the foundation of the ContentOps supply chain. Advanced ecosystems can have multiple content sources, but most initial supply chains will have a primary content source. If you do have multiple content sources, you'll need to have a plan for normalization, federation, or orchestration.

Create

ContentOps infrastructures rely on structured content, which means that they rely on structured content authoring tools. These tools can be simple web-form-style inputs, which are common for graph-based/digital modeling infrastructures. Or they can be more sophisticated environments for generating XML (eXtensible Markup Language) content.

The *create* stage can be an integrated part of the content management system that stores the source content. Alternatively, it can be composed of completely separate standalone tools, which are connected via API or workflow. Integrated authoring environments have become very common components of a web content management system (CMS). However, with enterprise-wide systems, it's common that one authoring environment won't be sufficient for all users, so implementations often use a combination of integrated and standalone authoring tools.

Regardless of the interface used to create the content, the most important consideration is that it natively creates content in adherence to the information architecture.

Manage

If the content source is the foundation of the ContentOps supply chain, the CMS is the bedrock of that foundation. The CMS needs to do much more than just store the content, though. Among those considerations are the following:

- Storage
- History
- Links
- Metadata/content typing
- Access
- Workflow
- Governance
- Tracking/analytics/reporting

Every system is going to implement these things differently. Since this functionality is generally considered table stakes for any modern ContentOps system, it's unlikely that any of these capabilities will be entirely missing from a chosen system. Instead, it's more a matter of validating that these functions work in a way that accomplishes your organization's goals. This is one of the places where a request for proposals (RFP) process can be dangerous. If you simply view these as checkboxes, every vendor will check them. This doesn't mean that every vendor will satisfactorily serve your needs. Be sure that you understand the true behavior of the CMS functions before finalizing your decision.

As a side note, if you're choosing an XML approach, the ISO 26531 standard—Systems and software engineering: Content management for product life cycle, user and service management documentation—dictates the most important functionality. It's worth checking out.

Connect

Modern ContentOps ecosystems are continually evolving into multiplatform software implementations. It's critical that the key technology components have the right APIs to support current and future components of the ecosystem. APIs become the basis of the apps and integrations that your team builds in the future, so in many ways, the total potential of a content ecosystem is the sum of its available APIs. The content source must have APIs available that allow access to the core features it provides around content management. Beyond that, the necessity for other APIs depends on the needs of your organization.

Not all of these APIs need to be provided by the content source; in fact, many of them, especially read-only, will be built on the content services. At its core, that's really what I see as content services—it's an API and surrounding utilities and functions.

As of this writing, most APIs are either RESTful or GraphQL. A comparison between these two types of APIs is beyond the scope of this book. Suffice it to say that both have their advantages; your organization should evaluate them against their likely usages.

Content Services

Content services come in a wide range of configurations. On the simplest end of the spectrum, you might have a PDF that is distributed to users. A more modern implementation would be a well-documented API that provides

appropriately permissioned access to content and associated metadata as JSON and HTML. The reality is that most ContentOps implementations will have a variety of content supply chains that may or may not leverage multiple content services layers.

Delivery Apps

Anything that a user uses to find, read, or otherwise access content is a delivery app. Mature ContentOps implementations will have many delivery apps. In fact, some of those apps will likely be the products the content is supporting.

Avoid creating substantive content in your delivery systems and apps. One of the most common examples of this is when groups create content in a learning management system (LMS). This is probably also one of the most forgivable, but it still carries all the problems of point system authoring. This content becomes an island. At best, it will be inconsistent with content created in other systems; at worst, it will be inaccurate or out of date.

Source Content Format

In your ContentOps technologies, the content itself has technology impacts. If you take nothing else away from this chapter, it should be that the choice (or your lived, operational reality) of content formats is the most consequential choice an organization can make when implementing a ContentOps ecosystem. This choice is typically made alongside other information architecture decisions, but as the source format is directly connected to the content systems you implement, it's good to look at them from a technical perspective here.

All content formats have advantages and drawbacks. The more rigid your structure, the more predictable it is for users and developers—and the more challenging it is to accommodate unforeseen cases. The more semantic your content, the more future ready and searchable it will be, but it will also be more challenging for infrequent authors to contribute to.

There are really four dominant approaches to the actual source content format:

- Proprietary formats.
- Text-based formats, predominantly Markdown and reStructuredText.
- Graph-based, generally via a headless system.
- Tree-based, generally through XML standards, predominantly the Darwin Information Typing Architecture (DITA).

HTML—Special Note

You'll notice that I didn't list HTML here as a separate content format. This is a debatable position. In my view, HTML isn't a source content format; though it is often used as the basis for proprietary formats, it's really a presentation format. My position is that when HTML is used as the source content format, it is either constrained to meet the needs of the tool or system, thereby effectively becoming a proprietary XML format, or its use in this way doesn't align with the goals of ContentOps, because the author is creating single-channel content in the presentation format.

Proprietary Formats

The one piece of universal advice that I feel confident providing is to avoid proprietary formats. Proprietary formats will lock you into a vendor's toolset and network, which creates a lot of risk and almost always ends with a painful and expensive conversion at some point in the future. Additionally, be wary of vendors that claim export compatibility with open standards as a substitute for true compliance with a standard. This can be tricky to suss out, so it's often good to engage a knowledgeable third party to help with this process.

Text-Based Formats

Text-based formats are another interesting case. They come in a lot of different flavors and can often be sub-components of other content formats. For example, they are often integrated into digital modeling. But some viable ContentOps systems do use them as standalone content objects that support the entire use case.

In the case where the ContentOps system is designed to work exclusively with a text-based format, this is typically a docs-as-code (Holscher, n.d.) approach to ContentOps. We're not going to delve too deeply into the usage or implementation of these systems here for a few reasons. First, these systems tend to be either very simple and single-purpose, or very complex and homegrown. The single-purpose version of these systems can be effective for a group's immediate needs and can be fairly classified as ContentOps implementations, but since this rarely scales up to support the broader goals of ContentOps, it's not really in the scope of this book. The very complex and homegrown systems aren't something we recommend, so we won't be spending time there either.

Graph-Based Content

The term “graph-based content” isn’t widely used, but it’s the most implementation-agnostic way of describing content based on an information architecture where content is flat, discrete units that connect based on typed relationships. Most often these are the result of digital modeling, which provides the content blueprint for describing other objects. One simple example might be a concert venue modeling a show. The show would connect to tickets and performers, which may then connect to other things, thus forming a graph.

Tree-Based Content (XML)

Tree-based content is really just another way of saying XML. XML is based on tags and allows for nesting content structure and semantics, and this nesting creates trees. XML is widely used in many content and data applications—web pages being a great example. HTML is just a defined subset of XML. As it relates to ContentOps, XML is an ideal format for technical and knowledge-style content.

Digital Modeling

Digital modeling of content is the practice of creating semantic blueprints for content structures that specifically describe your business or function. Typically, the digital modeling approach will result in a graph-style structure of content.

Graph-Based vs Tree-Based

To be clear, graph-based content and tree-based content are not mutually exclusive. Linking in XML can form complex relationships, and those relationships can form graphs. As a matter of fact, they must form graphs of some sort. The real difference here is that linking in XML tends to be less semantic, and therefore the relationships between linked objects have less meaning (think about an `<a>` tag or `@href` in HTML). This isn’t necessarily a good or bad thing; it makes XML more flexible for standard document, component, or snippet-style content. So XML can also be the basis of graph-based content, but in practice it typically isn’t exclusively graph-based in the way of content built on an information architecture resulting from digital modeling.

Graph-based tends to work well for content that is primarily accessed on the basis of its relationships, but less well for content that is inherently ordered or human-organized. Said another way, if your content lends itself to having a table of contents or logically organized sections, the pure graph-based approach may not best suit your needs. Inversely, if you need the maximum flexibility and programmatic assembly, graph-based is a very powerful structure.

Implementing a ContentOps System

Vendors

Any organization implementing a ContentOps ecosystem is almost certainly going to be working with one or more technology vendors. Bringing vendors into your process at the right time is critical. When you're making a change this large, getting vendor input early enough that you're still flexible in your planning can be a huge help. A major pitfall that many organizations run into, especially when running internal request for proposal selections, is designing criteria to match their past workflows rather than evaluating whether new technology can fundamentally change their workflows. The risk this needs to be balanced with is that vendors can also push an organization to do it their way when this isn't the best approach for the organization. Balancing the benefit of input from vendors with the risk it brings is key to getting the most from your relationships with them. Any organization implementing a ContentOps ecosystem is almost certainly going to be working with one or more technology vendors.

Proven Process

Start with the vision of where you think you want to be eventually. It will be a reference point as you work through your implementation, but not the initial focus. Move backwards to the 1.0 goal and create a clear definition.

1. Fill out the battle card below.
2. Create a plan
 - a. Scope
 - b. Information architecture
 - c. Technology blueprint
3. Implement
 - a. Choose technology
 - b. Migrate content
 - c. Stand-up and integrate technology
 - d. Configure output and UX
 - e. Train users

Filling in the Battle Card

The key questions we're going to answer in the process of filling in this battle card are: Where? What? Who? Why?

ContentOps Implementation Battle Card			
Channels	•	Content	•
Collaborators / Constituents		Key Business Objectives	
•		•	
Major Challenges		KPIs	
•		•	

Channels

Pretty simple. List the places you need your content to end up in the foreseeable future. This doesn't need to be complex for the battle card.

Content

What shared rules (metadata?) will all content be forced to have to be a valid part of the ecosystem?

Start by cataloging the content that will flow through the ecosystem. This could be a fairly short list, but considering potential future sources can be time well spent.

For each type of content, answer the following questions:

- Where does this content currently exist?
- Should it stay there?
- What is the state of the current content?
- Is it currently sufficiently structured?
- Does it currently have enough metadata to meet business objectives?
- What is the minimum number of content source systems that could contain this content?

- What special properties does this content have?
- What upgrades are required to make this content a viable input into the ecosystem?

Collaborators/Constituents

Catalog who will be involved in the project and who will be benefiting from it. If the list is long, consider creating a separate document that describes each of these groups in more detail.

Think through what each group may care about that doesn't appear in the key business objectives. If you're unsure, it's worth short conversations with each group to make sure they're aligned. Change management and consensus are always the biggest threats to successful ContentOps implementations, so skipping over this step can damage an otherwise well-planned and executed implementation.

As you're going through this step, it's worth remembering that for some groups, the highest priority is having to tolerate as little change as possible—and a goal here is to avoid annoying the content producers. Groups that must interact with the system but are not direct beneficiaries of the gains that it's providing have little motivation to support change. Contributors, reviewers, and other infrequent users often fall into this category. Working to understand their priorities and tolerance for learning new tools and methods is an important step in the process to a successful implementation.

Key Business Objectives

Objectives are not KPIs—they are the reasons why your KPIs matter. What are you trying to accomplish? The degree to which you're trying to accomplish it can be listed in the KPI section. Keep this conceptual. And keep the list short.

This list is the part all key decision makers need to agree on. If there isn't agreement that these objectives are worth the resources you're devoting to the project, then it will never get off the ground—or worse, it will sputter once it encounters the first major challenge. This sounds like common sense, but more than a decade of projects has taught me that it's all too common for business objectives to be not clearly defined, or sometimes not even considered.

Major Challenges

List all of the significant items or events that are (or might be) standing in your way. They can be technical or human challenges. You'll never predict all of

them, of course, but brainstorming some of the options up front will make you far better prepared for the curveballs.

KPIs

Every organization is going to have the big three to seven things that it cares about, and each of these should be captured with a KPI. As variables, KPIs are best when they're quantitative, as these are typically the easiest to communicate across teams. When it's not possible to form a quantitative KPI, these can be qualitative and/or binary. Binary is better than qualitative—at least for accurate measuring and benchmarking purposes. Binary KPIs can be statements like the following:

- “All of our content will now have a consistent look and feel.”
- “We’ll have one search that covers all of ...”
- “Now we can publish content to ...”
- “We’ll now have a governance that puts us in compliance with ...”

Each of these statements certainly has a quantitative impact hidden behind it, but it's often the case that content teams haven't been managing the “before state” well enough to make getting to a quantifiable improvement realistic. When you run into this case, it's best to establish metrics and measurements for future improvement, even if you're going to accept a binary KPI for your current project.

Plan

The battle card should present a good blueprint to begin planning the ContentOps tech implementation. From here, the structure of the plan can be based on the style and needs of your organization. The three sections I'd strongly encourage you to include are:

1. Scope
2. Information architecture
3. Technology blueprint

With these three sections well defined, there is a single place in which anyone involved can access the key details.

Implement

The steps in the process of implementing your ContentOps ecosystem have a general order, but many of them will happen in parallel.

Choose Technology

Technology selection will be a progressive process. Typically, ContentOps ecosystems will have more than one technology to acquire, but in some cases they may be built around a single system—at least initially. In almost all cases, there will be a product or app that is the driving technology, and it will dictate decisions on the rest of the stack because each following piece will need to integrate with it. While it's most common (and generally best) for organizations that are starting from scratch to choose the content source system as the driving technology, any of the three primary pillars can be your driving technology.

If the content source is your driving technology, it's easy to go from the information architecture decisions made in the planning phases, since most content source systems have a strong preference for which information architectures they'll support. If you're starting with content services or a key app, then you'll need to be sure that you have a clear path to connect those back to a content source that supports your information architecture.

When starting with a key app, such as the digital experience platform that has been implemented by the organization, it's important to avoid tunnel vision by placing too much focus on this one component. Implementations that are entirely driven by a need to supply a single major app can quickly be tied into knots on the back end and end up with an implementation that falls short on future requirements.

Migrate and Upgrade Content

Most ContentOps implementations for organizations starting from unstructured sources or proprietary point systems will need to do some form of content migration. Content migration is typically some combination of automation and manual conversion. The amount of manual labor required is generally determined by the complexity and scale of the migration. Complexity is driven by the state of the source content and the required structure of migrated content.

Low-complexity conversion can sometimes be entirely or almost entirely automated. Nevertheless, the conversion should be approached with extreme caution, because it's very common for the perceived complexity to be lower than the actual complexity. A mostly automated approach is always going to produce a lower fidelity result than one that includes more manual effort. There are situations where this fits inside of the business objectives, but often a mostly automated approach ends with an unexpected post-conversion manual effort, which can be especially challenging because it's unplanned and is harder to outsource. As such, it can fall to internal teams that are far more expensive than outsourced options.

It should be noted that standing up the technology is a step listed after content migration. This won't always be the case, and even when this reflects

the starting point for each, they'll run in parallel. That said, I have ordered them in this way because when establishing a ContentOps ecosystem built on a widely supported open standard, this step should be considered separately from the tools that will support it. Separate doesn't mean isolated, of course. The systems will always impact decisions in the content migration. When possible, however, my recommendation is to keep the content migration based on the information architecture.

Stand-Up and Integrate Technology

Once technology is selected and there is a clear path to moving content into it, it's time for the fun to really begin. In the modern era of SaaS (software as a service) systems forming the bulk of our content tools, turning on the key components of the system is generally a quick process. The complexity arises from the integrations between the systems.

Often, the chosen systems will have preexisting integrations. You shouldn't assume that these will do exactly what you need. Like anything else with software, this simply becomes a matter of discovery and iteration. The most important lesson to keep in mind is that some resources should be allocated to accomplish your goals. In many cases, this is an ongoing effort, and this is when organizations must be committed to persistent improvement in content experience. For these living ContentOps ecosystems, the initial stand-up is the starting line—and not a one-time event.

Upskill and Train Users

We all know that people are more important than technology, and success depends on prioritizing *people* upgrade as much as technology upgrade. Think of it this way: Do you have any doubt that a skilled contractor with a hammer would build a far better house than the average person with the most expensive nail gun on the market?

Training is almost always a function contracted with third-party providers. If you've built your ContentOps ecosystem on open standards, you should make sure to invest in tool-agnostic (but still aware of—and using—your tools) training on the concepts of that standard. Once key people understand the actual content and how they fit into the content strategy, getting tool training for daily users is a must.

Finally, the most often overlooked training is general training on the whole ecosystem. This is often a process that you build internally or in collaboration with a consultant who probably has been helping through the process. Having the ability to take someone through a high-level training on the ecosystem gives you a consistent starting point for everyone who is involved, no matter the person's

specific role. This provides a common understanding, and when done properly, a common lexicon. It's difficult to overstate the value of this education stage.

Configure Output/UX

Some content will make its way to users through integration with apps. The rest of the content will be deployed through systems that are primarily presentation systems for the content.

The look, feel, and function of the final presentation content is ultimately the most important result of your ContentOps implementation. The possible number of places and presentations of content is staggering and beyond the scope of this chapter. The key concept I'd like to impart here is that you should work from source to output, not the other way around.

A Note on GPT and AI for Content

As the first edition of this book was being prepared for publication, ChatGPT was released to the public, and overnight the landscape of content technologies changed in a major way. This is an extremely fast-moving technology, so almost anything written in a book, beyond basic explanations of its fundamental building blocks, will likely be out of date by the time it's read. Still, the impact on content professionals is too great to ignore in a book about Content Operations, so I will attempt to lay out some initial thoughts here.

As of this writing in June 2023, GPT (generative pre-trained transformer) AI (artificial intelligence) has gone through an exploration that has revealed many interesting and surprising strengths and some significant weaknesses. So with the understanding that we don't know what the terminal capabilities of this round of AI technologies will be, here is where it presently appears GPT/AI systems perform well with content related applications:

- Creation of tasks when the output is highly structured and many examples are widely available. Programming in popular languages and creating DITA XML structures are examples of this.
- Creation of tasks where accuracy of the output is unimportant and the goal is generating a “creative” work or a starting point. Image generators and presentation creators are both great examples here, but anything that comes of ChatGPT also fits this categorization.
- Summarization of “medium”-sized, discrete chunks of content. In practice, this means you can feed an AI model a dozen or so pages of content, ask it a question, and it will generally provide a good answer derived from that content. This contrasts with asking a question to a large language model

without giving it the content to pull from. In this case, the current GPT systems (ChatGPT 4 being the leader at the time of this writing) are not highly reliable. This problem is commonly referred to as hallucination.

- Anything prediction-based. Especially when purposefully built (trained), current AI systems seem to be very good at pattern matching and sorting. For content professionals, the most immediate use that comes to mind is metadata enrichment of content objects.

Additionally, it's almost certain that there are several narrow automation applications for AI, which will be organization specific. AI seems to be very good at single-step functions in domains that have well-structured rules with plenty of examples in the public domain. Ask it to do something discrete and it will do it if it's within its domain and span of control. The best way to think about this aspect is to imagine that you have the world's fastest and most knowledgeable assistant available to you, but this assistant has no grasp of strategy or multistep thinking. When you run into manual work that involves repetitive, complex, single-step operations, this may be a good candidate for AI.

Conclusion

While ContentOps is not technology, technology is both the enabler and the limiter of a Content Operations practice. Since a formalized ContentOps team and process will be organized around the capabilities of the tools in use, technology choices are a fundamental part of implementing ContentOps. Teams need to find a balance between designing a theoretical ContentOps practice that doesn't consider tools and designing one specifically around tools.

References

Holscher, E. (n.d.). *Docs as code*. Write the Docs. Retrieved June 17, 2022, from <https://www.writethedocs.org/guide/docs-as-code/>.

Indeed Editorial Team (2021). *What are battle cards? (and how to use them in sales)*. Indeed.com. Retrieved June 17, 2022, from <https://www.indeed.com/career-advice/career-development/battle-cards>.

Epilogue

You're Here because of People

Jonathan McFadden

When Carlos Evia asked me to pen this epilogue and take a strong position on the interplay between content and diversity, equity, and inclusion (DEI), I wanted to curl up in a corner and cry. As I sat at my keyboard, I realized I had perhaps said “yes” too quickly.

I’m not a Content Operations professional, and I’m not an academic writer. I’m a content designer whose foray into this industry came by way of journalism, marketing, and content strategy. I’m a relative newcomer. What do I know about the intersection between ContentOps and DEI? All I know is that I’m a DEI loudmouth who’s delivered a few presentations, written a few articles, spoken on a handful of podcasts, and helped found an employee-led DEI group at a company that I’m now two jobs removed from.

I don’t have any certifications. I don’t lead workshops. I don’t consult. I’m no expert. I just talk about it because I care about it.

And then it hit me: I care about it, which is precisely why I needed to write this. While I don’t possess the vaunted pedigree of the authors whose thought leadership fills this text, I do possess a quality that makes me uniquely positioned to offer commentary on this crucial topic. I’m a human whose life experiences have been shaped, influenced, and impacted by harmful, offensive content.

How to cite this book chapter:

McFadden, J. 2024. Epilogue: You're Here because of People. In: Evia, C. (ed.) *Content Operations from Start to Scale: Perspectives from Industry Experts*. Pp. 147–154. Blacksburg: Virginia Tech Publishing. DOI: https://doi.org/10.21061/content_operations_evia_11. License: CC BY

I'm willing to bet it's the same for you, too. Why? Because harmful content exists in way too many places: apps, websites, billboards, TV shows, signage, building design, your neighborhood, your job. It's so ubiquitous that it's nearly unnoticeable. And it's nearly unnoticeable because we're just *now* entering an era where more people are beginning to identify and unpack the potentially damaging effects the content they create can have on the people who digest it.

When I'm doing my job well (which, I hope, is more often than not), my thoughts don't dwell on the *what* as much as the *who*. Who is experiencing this? What are they feeling? What are they doing while they experience it?

Developing strategies and methodologies for operationalizing the way we make content is important. But it's not the reason you've come here—not the only reason, anyway.

People. You're here because of people.

More specifically, you're here for the individuals who consume, use, and interact with the content you create to accomplish tasks, meet goals, and fulfill obligations. They're the lifeblood that drives content creation in the first place.

Like the content we make, these people exist and experience the world differently, living their lives and filling space within an array of social, cultural, political, and physiological contexts. And our ambition to engineer and execute on content that resonates with as many people as possible, in spite of their aforementioned variances, relies on our ability to tap into these contexts, understand them, and deliver information in a way that helps, motivates, and empowers.

This is an idealistic and perhaps saccharine way of thinking about a job designed to serve business goals and generate revenue (I'm the first to confess that I'm preachy). Still, even though our work often happens in the confines of corporate settings, I don't believe our content functionally exists within enclosures. It leaps from digital devices to inhabit the micro-moments of our end users' daily habits. It builds digital communities where real life is constantly happening. And for better or worse, it fuels experiences—good and bad.

At its best, content that keeps inclusivity in mind amplifies trust and fosters belonging. It builds bridges, creates connections, and respects its users. At its worst, it diminishes identity, perpetuates systemic harm, and incites cruelty. The worst is what we want to avoid.

That's why this epilogue focuses less on processes and procedures and more on the inclusive content principles that should undergird them. Because we create for people, we *must* approach our work with a laser-sharp focus on DEI, which concerns itself with the representation and presence of *people* who differ in regard to age, race, gender, ability, sexual orientation, language, geography, and more.

These are the people who consume our organizations' content. They're not monolithic. They don't feel or think or believe or exist the same way we do. Yet we make content for *them*. And if we're not conscientious and careful, we'll

make stuff that adds to the orchestra blasting an all-too-familiar refrain to marginalized people: You don't belong here.

People First, Content Second

Because content doesn't create itself, I consider it a reflection of the people who create it and a testament to their choices during that creation process. That's why it makes such an indelible impact on the people exposed to it.

Content is created by and focused on people. That means it always has, and always will, find its place at the nexus of diversity, equity, and inclusion. It's inherently a DEI matter because content is inherently a people matter. The two don't merely exist in parallel. They are inextricably linked. The good news is that many of us who *do* content are paying attention and catching up.

Entire books have been written about implicit bias and the role it plays in product design. I strongly recommend you purchase *Design for Cognitive Bias* by David Dylan Thomas (2020), who explores how we can use those pesky mental shortcuts to craft more inclusive experiences for our users.

Just like a major tech company's facial recognition algorithm can offensively misidentify darker-skinned people as gorillas (BBC News, 2021), content practitioners' failure (or worse, unwillingness) to acknowledge, own, and confront their biases *will* inevitably show up in whatever content they create. These misfires don't happen inadvertently. They're historic and systemic—the result of a team, organization, or person's failure to prioritize inclusivity as a mode of working in content development.

Creating inclusive content represents a commitment to creating humane, responsible, equitable experiences. To borrow from Webflow's blog, it means "providing a safer experience for more people. It's not a badge of honor you earn with a single initiative or intention—it's an ongoing and intentional effort" (Kirchner, 2021).

Oftentimes, conversations about inclusive content start with the content itself. We find chatter about the importance of inclusivity and accessibility on LinkedIn posts and conference stages. Well-established content teams at some of the world's largest technology companies have created guidance around inclusive language in the form of style guides or addenda to existing writing manuals.

I praise these efforts. They're important. Necessary. But they're only a fraction of the real work that needs to happen to ensure that creating equitable and inclusive content becomes second nature.

We don't need to start with the content. We need to start with the people who make it.

In the corporate landscape, where a lot of content work happens, intentional efforts to focus on people have gained considerable traction. Over the last few

years, corporations the world over have started to recognize the value of integrating diverse and inclusive practices into their business models, drawing a direct correlation between diversity and inclusion and increased revenues.

An ample amount of scholarship exists to offer evidence on DEI's importance in the workplace. Some notable research, including studies authored by McKinsey & Company, shows that companies making a deliberate investment in representation at leadership levels experience a higher likelihood of financial outperformance. One oft-cited McKinsey study, published in 2020, shows that companies in the top quartile for gender diversity on their executive teams were 25% more likely to enjoy higher-than-average profitability than companies in the bottom quartile (Dixon-Fyle et al., 2020). It goes on to show that companies with ethnically diverse leadership teams reported a 36% uptick in financial performance over companies in the bottom quartile.

For companies looking to attract and retain top talent, DEI has become a vital part of the employee experience. A 2021 report from job review website Glassdoor shows that 76% of job seekers and current employees consider a diverse workforce an important factor when weighing companies and job offers (Glassdoor Team, 2021). Nearly half of those surveyed said they wouldn't apply to a company if there's a lack of diversity among its workforce: those numbers were especially high among Black and LGBTQ job seekers (41% for both groups).

How do these endeavors shape our content teams, which, anecdotally, are still mostly white? Actual demographic data for content strategists and content designers are hard to come by, but according to the 2016 Design Census by AIGA (the professional association for design), 73% of the surveyed designers were white (Howarth, 2017). In 2021, AIGA identified addressing and advancing diversity, equity, inclusion, and accessibility efforts as an ongoing opportunity area in the design community (AIGA, n.d.).

If we're to believe that the demographic makeup of an organization's leadership team has a direct influence on its financial output, shouldn't we also assume that the demographic makeup of a content team has a direct influence on the content its people create?

In the wake of the 2020 murder of George Floyd, I observed a shift in workplace emphasis on DEI. This visceral, videotaped instance of racialized violence riled the globe, forcing the United States in particular to confront and reckon with its history of racism, widespread inequity, social injustice, and white supremacy. As a result, company-sponsored equity workshops and diversity surveys mushroomed into candid conversations about the perils and ills of white supremacy and systemic racism. Corporations frantically raced to release their diversity numbers, and corporate communication teams authored beautifully crafted statements that live in perpetuity on their social media pages.

That same tide of introspection swept over the content community, bringing a flood of conversations, presentations, books, and blogs about inclusive

content and methods for expunging racist, ableist, ageist, sexist, homophobic, and other harmful language, components, and biases from products their companies created. Perhaps with more stark clarity than ever before, content practitioners working on products, websites, blogs, videos, social media, and the like began to unpack the ripple effect the content they create has on the people who absorb it.

I don't mean to suggest that this is a *new* conversation. Content folks were talking about inclusivity and equity in the industry long before George Floyd's murder. But as 2020's pandemic-borne restrictions forced many of us to sit with our feelings and confront the ugly truths about racism, those discussions elevated to enterprise levels, placing a lofty responsibility into the hands of everyone who creates content.

Now, a few years removed from George Floyd's death, has our content become more inclusive? Maybe. There are certainly more inclusive language guidelines. But how do our teams look? Can we affirmatively report that content has become more inclusive when content teams and content leadership, by and large, remain relatively homogenous? I'm not so sure.

Diversifying our content teams reduces homogeneity and offers the benefit of fresh perspectives, new ideas, different skills, and more engaged teammates. All those ingredients combined make for a potent mix that will invariably shape content strategy, content development, and the people to whom we serve content experiences.

Creating and advocating for inclusive content should be part of every content job description, just as it should be baked into any framework we use in the way we work. This philosophy extends to our talent as well. Content creators should represent a diverse cross section of backgrounds, experiences, abilities, languages, and thinking styles, for example. Diverse content teams make diverse content, bettering our chances of connecting with and appropriately serving the vast array of people using our products.

What's at Stake if We Don't Act?

By now, you may have realized that content doesn't exist in a vacuum. It's a living, breathing organism with real-life implications.

Content sends powerful psychological signals that help people understand whether the experience they're presented with is safe and welcoming or dangerous and threatening.

Those signals can be stark and clear, or they can be subtle and covert. Either way, they exist.

- We find them in the use of the words "master" and "slave" to describe the interdependent relationship between files.

- We find them in button text and acronyms that are incompatible with screen readers.
- We find them in résumé-scanning algorithms that show strong bias against women.
- We find them in hurtful, dehumanizing words like “alien,” “invalid,” or “other” appearing in online forms in reference to someone’s citizenship, name, or gender identity.

In 2021, I partnered with Art Schwartz-Restrepo, a fellow senior UX content designer, to deliver a presentation at Confab, a leading content strategy conference. We discussed ways we could use content to advocate for underrepresented users, or rather, people who have been historically harmed and marginalized by institutions and systems that diminish their personhood. Our goal was to demonstrate that people working in content—no matter where in the life cycle they begin exercising their influence—possess the capability to engender equity and humanity in products.

We used a case study from Stanford and Cornell Universities as a linchpin for our argument. In that study, researchers found that content created for a series of Facebook ads urging signups for online STEM (science, technology, engineering, and mathematics) courses inextricably affected the viewers’ ambient belonging, a psychological term referring to the feeling people experience when they feel like they are connected to and accepted by a particular group or community. The study found that when gender-inclusive verbal and visual cues were paired in the ads (for example, with images depicting a variety of women as students and ad copy assuring the audience that women were part of the target demographic), course enrollment among women increased by 18% without reducing engagement among men. However, when those cues were removed and less-inclusive imagery and text were rendered instead, such as a stock image of a laptop displaying binary code, the number of women enrollees dropped by 8% (Kizilcec & Saltarelli, 2019).

When Art and I started work on our presentation, we discussed the psychological impacts signs like “Whites Only” or “Colored Restrooms” must’ve had on Black people living in the Jim Crow-era south, when segregation and discrimination were codified as the “separate but equal” lie. Our contention was that while the meaning of those words—of that content—was evident in physical spaces, the words likely had farther-reaching psychological ramifications that couldn’t be detected visually.

That argument didn’t make it into the presentation, but I still believe it holds water. Content has long-lasting impacts on the psyche, even when the moment has passed.

I know this because I grew up with parents who were children in the sixties, near the tail end of the civil rights movement, but came of age as race relations in the United States slowly shifted. My father grew up in Lake City,

South Carolina, and my mother in Brooklyn, New York—two very different parts of the country racked by violence as integration grew normative.

And while some things *did* change, much didn’t. My dad remembers the restaurants that were only for white people and the looks he’d get when he dared to cross into their world. My mom recalls the streets and subway stops where classmates chased her because of her hair and skin. Although they’re decades removed from those incidents, the memories remain vivid and visceral.

Content we create may come and go, but it has permanence. It leaves impressions. It leaves feelings. And it can leave pain.

You Work for People

Designing and creating content for people requires a profound understanding of its power. These aren’t just words on a page. This isn’t just an interaction or animation. We’re not just building intricate user flows, e-carts, or checkout experiences.

We’re making space where real life happens, where real people form relationships, practice commerce, sign up for flu shots, and order food while they’re quarantined at home. What we create affects them in very real ways.

Content work, no matter its flavor, requires an empathetic bent and a deep understanding of the implications and consequences of that content. And content teams, no matter the industry in which they sit, need to reflect the range of intersectional identities, backgrounds, experiences, abilities, and languages of our users. An inclusive team is not foolproof. It’s not a fail-safe against creating harmful things, but it does yield benefits that homogenous teams don’t. Different ideas. Different perspectives. Different experiences. Together, all these fuel better experiences that are more human, more empathetic, and more inclusive.

Committing to very human ideals like diversity, equity, inclusion, accessibility, belonging, and representation isn’t simple. Such commitments require behavioral change, and change—even the right kind of change—is tough.

But it’s not impossible. We *can* do this if we care enough to do it. We can make content a more inclusive community and practice if enough people do more than endorse the idea—if they actively champion its continuity and success.

We’ve got work to do. Are you in?

References

- AIGA. (n.d.). *AIGA 2021 design POV research reports*. Retrieved June 20, 2022, from <https://www.aiga.org/aiga-design-pov-reports>.
- BBC News. (2021, September 6). *Facebook apology as AI labels black men “primates.”* Retrieved June 20, 2022, from <https://www.bbc.com/news/technology-58462511>.

Dixon-Fyle, S., Dolan, K., Hunt, V., & Prince, S. (2020, May 19). *Diversity wins: How inclusion matters*. McKinsey & Company. Retrieved June 20, 2022, from <https://www.mckinsey.com/featured-insights/diversity-and-inclusion/diversity-wins-how-inclusion-matters>.

Glassdoor Team. (2021, July 12). *What job seekers really think about your diversity and inclusion stats*. Glassdoor for Employers. Retrieved June 20, 2022, from <https://www.glassdoor.com/employers/blog/diversity/>.

Howarth, D. (2017, January 31). Design industry is 73 per cent white, finds AIGA survey. *Dezeen*. Retrieved June 20, 2022, from <https://www.dezeen.com/2017/01/31/design-industry-73-per-cent-white-lacks-diversity-finds-aiga-census-survey/>.

Kirchner, R. (2021, October 13). *How to create inclusive content that welcomes a wider audience*. Webflow. Retrieved June 20, 2022, from <https://webflow.com/blog/inclusive-content>.

Kizilcec, R.F., & Saltarelli, A.J. (2019). Psychologically inclusive design: Cues impact women's participation in STEM education. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (pp. 1-10). New York: Association for Computing Machinery.

Thomas, D.D. (2020). *Design for cognitive bias*. New York: A Book Apart.

Afterword

Jason Swarts

This was a good idea for a book. The essays collected in this volume all speak with wisdom and clarity about the issues and practices of ContentOps, and the overall message is especially meaningful to me, as an academic who trains technical communicators. As Carlos Evia mentions in his introduction, we are kindred spirits in academia. We both belong to academic programs that teach perspectives and practices that reflect, reinforce, or at least rhyme with the perspectives and approaches outlined in these chapters. We also believe that the academic field of technical communication is enhanced by stronger connection to the issues, concerns, and methods shaping the profession of technical communication today. Likewise, I would hazard to say that we also agree that research in academic technical communication can contribute to technical communication practice.

In fact, when considering what I could add in an afterword to an already detailed and comprehensive look at ContentOps, it seemed appropriate to focus on an argument describing how academia and industry benefit from reflecting one another, especially on issues related to this field. After providing a bit of history, I will discuss some of the authentic connections between academia and industry that the chapters in this collection suggest. Specifically, I will focus on both the practical skills and the training to which new technical communicators should be exposed; I will also talk about some of the basic research going on in the field that can influence how we think about issues related to

How to cite this book chapter:

Swarts, J. 2024. Afterword. In: Evia, C. (ed.) *Content Operations from Start to Scale: Perspectives from Industry Experts*. Pp. 155–162. Blacksburg: Virginia Tech Publishing. DOI: https://doi.org/10.21061/content_operations_evia_12. License: CC BY

ContentOps. Following that discussion, I will lay out what I take to be the educational mission that derives from the professional practices discussed here.

Two Cultures?

Let's start with the big issue: the "divide" between technical communication as a field of study (i.e., academia) and technical communication as a professional practice (i.e., industry). Although I have hung scare quotes around the word "divide," there is no denying that the gulf between what academics write about and what practicing technical communicators want to read appears wide. This observation does not even address the point that academic journals are hardly a forum that attracts the attention of academics and practitioners equally. Even industry-focused journals like *Technical Communication* have trouble meeting that expectation (see Andersen & Hackos, 2021). The divide appears real and not unlike the "two cultures" divide that C.P. Snow saw in science and the humanities (1959).

The problem with a "two cultures" view of the issue, however, is that it promotes a deficit model, supposing that one of the cultures lacks what the other has. It should not be that way—and it need not be that way. Collections like this one highlight both the places where there can be convergence between industry and academia and the specific proficiencies, issues, and challenges that can bridge the divide. Mike Albers argues, and I agree, that although we do not need to look hard to find mention of this divide, we "sincerely hope there are no future texts talking about the great split, as its occurrence would be a great loss for both sides" (2016, pp. 293–294).

For the sake of making a point, however, let's take a moment to consider aspects of this divide. One element to note is the difference in motivation that leads to differential consideration of topics that need attention. In industry settings, the need may be customer driven (see Kimball, 2015, p. 142) or driven by internal organizational needs. Often the need is problem-based and time-sensitive, focused on recurring problems that need solutions (see Johnson-Eilola & Selber, 2013). In academia, needs are not driven by the same motivations, but they could and should be. Trends in technical communication practice, like writing for reuse, for localization, for personalization, for findability, for searchability, and for Content-Ops in general, characterize the skills and capabilities of technical communicators that are needed in the field (see Andersen & Evia, 2019).

An academic field is not sustained, however, and cannot develop a center by jumping from one applied problem to another. Academic research that is not problem-oriented may instead be motivated by more basic needs to understand concepts like audience, context, location, networks, and collaboration—concepts that do, clearly, influence practice. Not all scientists, social scientists, or humanists of any variety devote themselves, strictly, to applied and instrumental

purposes. Basic research must happen in a field as well: knowing for the sake of knowing, answering questions that are unknown for applications that we have not yet imagined. But regardless of the kind of research done, there is a need for good communication, for academics to speak to industry partners and convey what they know that may be of help now or of interest and potential value at some later date (see Albers, 2016; Andersen & Hackos, 2021).

Lest I get too far away from the subject at hand, let me tie this thread into the tapestry. ContentOps, as expressed and developed across these essays, presents a field of practice and study that is rich with opportunity for both applied research and basic research, and there are authentic connections to be identified at the intersection of the two. The essays in this collection present opportunities to identify these authentic connections and to point out the research issues that are revealed at the interstices. The essays are also excellent starting points for identifying how to train new generations of technical communicators.

Authentic Connections

I would like to elaborate a few of these connections and write about what I see as areas of commonality, practical issues to be solved, and basic research to be done. Those areas are:

- Rhetoric and stylistics (understanding: operationalizing, localization).
- Publication management (understanding: ContentOps, governance, business).
- Experience design and audience analysis (understanding: customer experience, personalization, search support).

For years, I have taught a foundational, required seminar for our Master's of Science in Technical Communication program. It is a course on structured, topic-based authoring. Students learn how to write topics using discrete content elements and then structure those elements together into recognizable genre-specific content. Specifically, we learn to use DITA to build documentation projects. We learn about translation and localization and single sourcing. I never run out of subjects to cover and practice, but my students sometimes run out of patience. I mention this because I have learned that the value students see in a course like this is almost always retrospective. Some tune in right away, but others need to have the experience of doing an internship, applying for jobs, or working on a team of technical communicators to realize that there really *was* a need to learn how to write in topic formats, to pay careful attention to word choice, to include clues in topics to aid in searchability and findability, and to give attention to how to structure those topics to create a beneficial user experience.

Rhetoric and Stylistics

As the chapter on operationalizing ContentOps makes clear, this kind of writing and editing process is far removed from one in which a single author works in isolation on a single publication from start to finish. This craft model of writing is being replaced by team-based, incremental, distributed approaches to writing (see Brumberger & Lauer, 2015). Aside from learning a new model of writing, students also need to learn new ways of thinking about the kinds of rhetorical work that these texts do and how that work is supported by the stylistic choices made in those documents. For example, many students initially struggle with the concept of a topic (as opposed to a document). They are unused to writing content that can stand alone, relatively free of context, to be combined with other topics into various output formats.

Students can learn the practice of topic-based writing, of course. The tools supporting the management of topic-based writing and the review of topic-based writing can be learned as well. Students might groan under the weight of learning to use a structured authoring tool or learning to push and pull files from GitHub, but those practices can be modeled and taught just as easily. However, there is basic research to be done here as well. Consider, for example, what a topic even is. What are the characteristics of a topic-based style? What kind of rhetorical work must topics do to help engage readers, structure their use of the content, and lead them to additional content? Observational research can answer some of these questions, but the field also benefits from other approaches, such as corpus analysis, which attempts to look across hundreds or thousands of topics to find stylistic patterns that people doing ContentOps have tacitly understood to be important and have grown accustomed to using. With this basic research, we can discover how language choices anticipate and mitigate some of the problems that readers can encounter when using topics (see, e.g., Swarts, forthcoming).

Basic research on writing style and writing choices can also help with our understanding of ContentOps-like localization (e.g., Getto & Sun, 2017; Maylath & St. Amant, 2019). The sometimes subtle and not-so-subtle rhetorical and stylistic choices that writers make in presenting their content can make huge differences in the uptake and use of that information (Agboka, 2013). Close attention to stylistics and rhetorical process can also influence the efficiency of translation (Kohl, 2008) and reader engagement (Gonzales & Zantjer, 2015).

Publication Management

A related set of connections are highlighted in chapters on the practices of ContentOps, on governance in a content development process, and on the business side of ContentOps. In academic technical communication, these

issues show up in courses on publication management and content management. Although classrooms are poor substitutes for the organizational crucibles in which ContentOps is done in the wild, they are safe spaces where these practices can be modeled and made the object of discussion.

Basic research on these topics also has much to show us. In their integrative literature review on project management, Ben Lauren and Joanna Schreiber show how the literature focuses on topics like teamwork, collaboration, and interpersonal relationships (2018), but there is a need for additional basic research on the governance and the business side of ContentOps. What are the most productive kinds of working partnerships to cultivate for effective management of content and review processes (see Batova, 2019)? What are effective ways of managing content production and stakeholder engagement (see Sedmak et al., 2019)? Additional questions and projects are to be found within this collection.

Experience Design and Audience Analysis

There is also the connection between work on user experience and audience analysis as they connect to topics like understanding customer experience, personalization, and search support. Each of these ContentOps issues presents clear problems of application. How do we develop content that customers are going to find engaging and that they will stay invested in? Here, basic information on engagement (Khaslavksy & Shedroff, 1999) provides awareness of mechanisms through which engagement is built. In addition, research across domains of communication, such as comics (e.g., Yu, 2020) and video games (DeWinter & Moeller, 2016), provides lessons about how communication in these different media formats may employ techniques of audience engagement that can be adapted to the needs of ContentOps.

Personalization is also a concern that some of us have addressed in our teaching, attempting to understand how to implement practical solutions like developing a metadata strategy for tagging topics and serving up that content using conditional processing. But these coarse attempts at developing personalization can be enhanced further by looking at how (for example) algorithms and intelligent agents can process information and serve it to users (e.g., Ranade & Catá, 2021).

Basic research can also help us understand the kinds of experiential issues that arise when readers access information across a document landscape where content is served up on demand. Developing effective content for those kinds of environments depends on understanding search behaviors (Erickson, 2019; Pirolli, 2007) and how to respond with rhetorical choices that help texts do more of the work to help readers adapt the content to their local circumstances of use.

I do not mean for this afterword to give a complete overview of academic research in technical communication. In fact, the citations above only skim the surface and largely reflect my own selection of research that I find particularly useful in the context of trying to teach ContentOps.

An Educational Agenda

The problems and situations discussed throughout this book are real, exciting, and challenging. One implication of this conversation is that the technical communication workplace has been transformed by technologies and processes that support the complex work of ContentOps.

This collection of essays offers a thorough and useful perspective on technical communication work practices. All of the motivations and challenges laid out in these pages should influence how we talk about and implement the training of new technical communicators. But perhaps the most important thing accomplished in this collection is the very definition of a concept like ContentOps as an integrative description of communication and organization practices that reveal to us the ways in which our profession is evolving.

Technical communication educators like Carlos and myself have a responsibility to help prepare students to jump into these roles, not fully formed and fully capable, but rhetorically and technologically prepared to meet and understand the challenges. We have a responsibility to help our students comprehend the technologies of ContentOps. We also have a responsibility to help them understand how work, writing, and review practices are likewise changed by ContentOps.

At the same time, we can promote the intellectual history of technical communication as an interdisciplinary field and draw upon the methods and methodologies of social sciences, communication, user experience, usability, social cognition, and education to learn how to create and collect useful information about the contexts of ContentOps as well as about the efficacy of our current practices and the potential for new practices. We can determine how best to apply methods and methodologies like observational analysis, rhetorical analysis, and corpus analysis to determine effective ways of generating data about ContentOps and to understand how it works. The point underlying the chapters in this collection is that we need effective, reliable, and consistent ways of looking at our documentation practices and assessing whether and how they are working. We need research that leads to education.

These agenda items for technical communication research also point to a need to reorient technical communication education to focus on the articulation and application of deliberate design thinking (Tham et al., 2022) and to develop (and practice developing) heuristic ways of doing ContentOps to reach and engage audiences more effectively.

We must recall, too, that technical communication has humanistic roots (Ranade & Swarts, 2019). We have an obligation to our various publics to pay attention to user experience, ensuring that we give due and suitable attention to issues of diversity, equity, and inclusion (Jones, 2020). Users and readers remain at the center of the work that we do, even in light of the automation and efficiencies introduced by ContentOps.

References

Agboka, G.Y. (2013). Participatory localization: A social justice approach to navigating unenfranchised/disenfranchised cultural sites. *Technical Communication Quarterly*, 22(1), 28–49.

Albers, M.J. (2016). Improving research communication. *Technical Communication*, 63(4), 293–297.

Andersen, R., & Evia, C. (2019). What leaders in industry and academia think about technical communication training and education. *Proceedings of the 2019 STC technical communication summit*. Denver, CO: STC.

Andersen, R., & Hackos, J. (2021). Practicing technical communicators' experiences with and perspectives on academic publishing. *Technical Communication*, 68(3), 29–55.

Batova, T. (2019). Lost in content management: Constructing quality as a global technical communication metric. *Technical Communication*, 66(1), 30–52.

Brumberger, E., & Lauer, C. (2015). The evolution of technical communication: An analysis of industry job postings. *Technical Communication* 62(4), 224–243.

DeWinter, J., & Moeller, R.M. (2016). *Computer games and technical communication: Critical methods and applications at the intersection*. New York: Routledge.

Erickson, L. (2019). *The role of taxonomy and search in content usability*. [Master's thesis.] FH Joanneum. Retrieved June 20, 2022, from <https://epub.fh-joanneum.at/obvfhjhs/content/titleinfo/4403250/full.pdf>.

Getto, G., & Sun, H. (2017). Localizing user experience: Strategies, practices, and techniques for culturally sensitive design. *Technical Communication*, 64(2), 89–94.

Gonzales, L., & Zantjer, R. (2015). Translation as a user-localization practice. *Technical Communication*, 62(4), 271–284.

Johnson-Eilola, J., & Selber, S.A. (Eds.). (2013). *Solving problems in technical communication*. Chicago: University of Chicago Press.

Jones, N.N. (2020). Coalitional learning in the contact zones: Inclusion and narrative inquiry in technical communication and composition studies. *College English*, 82(5), 515–526.

Khaslavsky, J., & Shedroff, N. (1999). Understanding the seductive experience. *Communications of the ACM*, 42(5), 45–49.

Kimball, M.A. (2015). Training and education: Technical communication managers speak out. *Technical Communication*, 62(2), 135–145.

Kohl, J.R. (2008). *The global English style guide: Writing clear, translatable documentation for a global market*. Cary, NC : SAS Institute.

Maylath, B., & Amant, K.S. (2019). *Translation and localization: A guide for technical and professional communicators*. New York: Routledge.

Pirolli, P. (2007). *Information foraging theory: Adaptive interaction with information*. New York: Oxford University Press.

Ranade, N., & Catá, A. (2021). Intelligent algorithms: Evaluating the design of chatbots and search. *Technical Communication*, 68(2), 22–40.

Ranade, N., & Swarts, J. (2019). Humanistic communication in information centric workplaces. *Communication Design Quarterly*, 7(4), 17–31.

Sedmak, M., Ruggeri, K., Boldt, R., Dennis, S., & Forsythe, J. (2019). Visualizing content strategy for complex content ecosystems. *Technical Communication*, 66(2), 147–162.

Snow, C.P. (1959). *The two cultures and the scientific revolution*. New York: Cambridge University Press.

Tham, J., Howard, T., & Verhulsdonck, G. (2022). Extending design thinking, content strategy, and artificial intelligence into technical communication and user experience design programs: Further pedagogical implications. *Journal of Technical Writing and Communication*, 52(4). <https://doi.org/10.1177/00472816211072533>.

Yu, H. (2020). Conceptual art or readable contract: The use of comics in technical communication. *Technical Communication Quarterly*, 29(3), 222–239.

Contributors

Rahel Anne Bailie is an executive consultant at ICP, the person behind Content, Seriously, and a long-time content practitioner. She began her career in technical communications, eventually moving into content strategy, content operations, and now into creative operations, sorting out thorny content problems for clients who seriously want to use their content as business assets. Rahel is an elected STC and honorary ISTC Fellow, university instructor, workshop facilitator, certified UX practitioner, and accredited Cognitive Edge consultant. She is well known in the content industry, with awards, accolades, and books—most notably *Content Strategy: Connecting the Dots Between Business, Brand, and Benefits*. Rahel has spent the last few decades helping clients create order from their content chaos, developing and delivering a range of content courses, and writing about content very, very seriously.

Deane Barker is the global director of content management at Optimizely. He's been working in content management for 25 years and has written four books about the patterns and practices of modeling, creating, managing, and delivering digital content, including *Web Content Management: Systems, Features, and Best Practices* for O'Reilly Media and the just-published *Web Project Guide: From Spark to Launch and Beyond*. Deane lives in Sioux Falls, South Dakota.

Patrick Bosek is a co-founder and CEO of Heretto, an industry-leading SaaS solution for content operations and digital self-service. Since beginning

Heretto in 2005, Patrick has worked on a wide range of projects, all focused on improving the authoring, production, and distribution of content. He is a software industry professional specializing in developing, productizing, and solving problems with product content software. Patrick is a skilled developer, thoughtful manager, and passionate customer advocate.

Carlos Evia is a professor of communication and director of the Academy of Transdisciplinary Studies at Virginia Tech, where he is also associate dean for transdisciplinary initiatives and chief technology officer in the College of Liberal Arts and Human Sciences. A mediocre studio session DJ turned decent entertainment journalist, Carlos moved to the intersection of information technology and the humanities with early jobs as database designer and technical writer. In his academic career, he has been a professor of English (professional and technical writing) and communication (digital publishing and content strategy) and an award-winning researcher in transdisciplinary technical communication and content operations. He is the author of *Creating Intelligent Content with Lightweight DITA* and worked on the development of the Darwin Information Typing Architecture (DITA) and Lightweight DITA (LwDITA) standards for digital content. A proud Mexican American, Carlos taught himself how to speak English by watching *Sesame Street* and (not simultaneously) reading *MAD Magazine*.

Kate Kenyon is the head of content and CX operations at Chase International Bank. Over the past two decades, she has spent time in almost every role in content: writer, strategist, content modeler, and CMS wrangler. Her experience across content gave her the depth of understanding necessary to know what tools, processes, and environments creative teams need to scale—and how to build them. Kate now focuses on the enterprise-level design, set-up, and management of content and CX teams.

Jeffrey MacIntyre is a high-mileage, New York-based independent consultant. Over numerous tours of duty and two decades of swashbuckling through design, product, and data management settings, he has solved or struggled with every unruly content puzzle under the sun. He consults for Bucket Studio, focusing on “shovel-ready” solutions to personalized, automated, and simplified customer experiences. He runs Bucket Brigade, the only community for those who design connected experiences, and he writes Bucket List, a newsletter of tales from the trenches of his consultancy and the behavioral science behind simplification in design. His clients include Best Buy, Chase, Condé Nast, Consumer Reports, IBM, Mailchimp, Mastercard, NBCUniversal, News Corporation, Philips, Thomson Reuters, UNICEF, Universal Music, Viacom, Vimeo, and Yahoo! Jeffrey is a founding member of the Consortium of Personalization Professionals.

Jonathan McFadden is a senior content designer at Shopify, where he designs experiences for Shop, a consumer-facing shopping and order tracking app. A writer at heart with a penchant for product, Jonathan's career has included stints as a journalist, communications coordinator, copywriter, and content strategist. He's an unapologetically loud champion for diversity, equity, and inclusion, and he has delivered a number of talks at conferences and meetups centered on representation, inclusive writing, and the power of self-advocacy. In 2019, he helped co-found and co-chair the first diversity and inclusion group for one of his former employers and continues to hold space for conversations about race and equity for anyone willing to listen (and even those who aren't). When he's not designing with words as part of his day gig, Jonathan writes them for clients as the owner of Jon Writes, his content writing side hustle. He could tout a number of other career milestones, but his proudest by far is the day he survived a purse swing by a woman on trial for murder. He has no aspirations to surpass it.

Kevin P. Nichols is the executive director of experience for AvenueCX, a content strategy consultancy that helps businesses enhance and improve their content experiences. Prior to that, he grew one of the largest and most successful global content strategy teams in the world at SapientNitro and, later, at Publicis Sapient. In the mid-2000s, he worked at Molecular, with clients such as HP and Reebok, on content operations and customer-facing content experiences. In the early 2000s, he worked as one of the key architects for MIT OpenCourseWare, helping design and launch the overall publication process and content experience, and led the content production team. Kevin has a master's degree from Harvard Divinity School and a BA from Kenyon College. Currently, he resides in Medford, MA, and Cape Cod, and he is passionate about international human rights, particularly those of women and refugees.

Sarah O'Keefe is the CEO of Scriptorium. She founded Scriptorium in 1997 to work at the intersection of content, technology, and publishing. Today, the organization is known for expertise in solving business-critical content problems with a special focus on product and technical content. Sarah identifies and assesses new trends and their effects on the industry. Her analysis is widely followed on Scriptorium's blog and in other publications. As an experienced public speaker, she is in demand at conferences worldwide. Sarah holds a BA from Duke University and is bilingual in English and German.

Loy Searle has been a globalization and content industry leader for 20+ years. In the ERP industry, her teams pioneered single-sourcing content strategies and built extraordinary integrated global CMS/terminology solutions. At Google, she led Global Production and Language Services. Her focus was speed @ scale—turning language quality around while shortening time to market.

At Intuit, her team built a scalable globalization and innovative content foundation to support market expansion. At Deluxe, her team transformed their practice to support the entertainment industry's digital transformation. As past president and board member of Women in Localization, Loy is committed to the localization/globalization industry and the advancement of women within it. Today at Workday, Loy's team is building a Localization and Globalization CoE to scale and support the company's expansion goals.

Jason Swarts is a professor of English at North Carolina State University, specializing in the field of technical communication. He is a core faculty member in NC State's MS program in technical and scientific communication as well as in the communication, rhetoric, and digital media PhD program. Professor Swarts focuses his research on the technological mediation of writing practices, the rhetoric of technology, workplace communication, and emerging genres of technical communication. Previous work has appeared in *Technical Communication Quarterly*, the *Journal of Technical Writing and Communication*, the *Journal of Business and Technical Communication*, *Technical Communication*, and *Written Communication*. Over the years, his research has received a number of awards, including NCTE Best Article on Theory or Philosophy of Technical or Scientific Communication (2005, 2009, 2011), the Nell Ann Pickett Award for Best Article in *Technical Communication Quarterly* (2007), and the Frank R. Smith Distinguished Article Award in *Technical Communication* (2013). His first book, *Together with Technology* (Baywood/Sage), received the NCTE Award for Best Book in Technical or Scientific Communication (2009). He has also written *Wicked, Incomplete, and Uncertain* (2018) and, with Cheryl Geisler, completed *Coding Streams of Language* (2019).

Index

2020 B2B Buyer Behavior Study 87

A

- Abel, Scott 5
- academic field, technical
 - communication 156–157
- accessibility, content 38, 62–63
- Account-Based Marketing (ABM) 87, 100
- accounting process 18–19
- accounting transactions 18–19
- action-oriented approach 57
- adaptable content 55
- administrative metadata 46
- advertisements 9
 - breaking conventions 50
- advocacy 82
- affordances 67, 68, 69, 70, 71, 72
- Agile operating model 17

- AI (artificial intelligence) 144–145
- alphabetic syllabary languages 49
- An Uneven History of Content
 - Strategy (Day) 15
- analyzing customers 78
- Andersen, Rebekka 5, 6
- ANSI (the American National Standards Institute) 50
- approvals 37
- Atherton, Mike 4, 70
- audits, post-production 38
- authentic connections, technical communicators 157
- authoring and editing,
 - content 28–29
- automatic identification and data capture (AIDC)
 - technologies 51
- awareness 78, 81

B

Bailie, Rahel 33
 Barker, Deane 9, 13
 Batova, Tatiana 5, 6
 battle card 139
 business objectives 140
 challenges 140
 channels 139
 collaborators/constituents 140
 battle card sample 125
 binary approach,
 post-production 38
 Bosek, Patrick 67
 brand identity content 30, 35, 76,
 77, 78, 80, 81, 83. *See also*
 smarter brands
 advertising and marketing 81
 consideration, customer
 experience 81
 loyalty 82
 British Standards Institute 49
 bug tracking tool. *See query*
 management
 business assets 1
 business benefits 43
 business justification 26
 building business 32
 for consistency, content 30–31
 risk management and
 compliance 31
 scalability 26–27
 velocity 27–28
 business objectives 15, 126, 139,
 140, 142

C

Carroll, John 57
 CCMS (Component Content
 Management System) 69
 channels 126, 139
 ChatGPT 144, 145
 client and supplier,
 localization 112–113

clothing company, content
 governance 36
 cloud-based applications 114
 collaboration
 creating kinetic content 45
 operating models 17, 22
 collaborators 127
 collaborators/constituents 140
 Colman, Jonathon 83
 Confab 152
 confidence 30, 100, 105, 129
 configure output/UX 144
 consideration, customer
 experience 81
 consistency, content 30–31, 127.
 See also page consistency
 content 53, 139–140. *See also*
 specific contents
 commodification of 6, 7
 consistency 30–31
 definitions 4, 5, 6
 delivery 29–30
 maintenance 28
 operationalization of 8
 outlining and documenting 127
 personalization 10
 production 19
 rendering 29
 risk management and
 compliance 31
 robustness of 43
 terms related to 5
 velocity 27–28
 within an organization 21–22
Content Advantage, The (Clout 2.0)
 (Jones) 8
 Content as a Service (CaaS) 44,
 64–65
 content assets 19, 52, 53, 56, 94
 content delivery 64, 103, 130
 contextual 97
 dynamic content 65
 independent content 65
 offered content 65

- omnichannel 75
- profiled content 65
- service 64–65
- content development process 3, 25, 158
- content factory 25
- content governance. *See* governance
- content grammar accuracy 38
- content life cycle 4, 15, 22, 27, 28, 29, 31, 55, 75
- content machine 6, 7
- content management system (CMS) 68, 84–85, 93, 133
- content modules 58–59, 88
- Content Operations**
 - (ContentOps) 1, 13, 155
 - and customer journeys 87–88
 - architecting confidence 129
 - automate continuous delivery
 - pipelines 45–46
 - battle card sample 125
 - channels 126
 - collaborators 127
 - components of 130
 - consistency 127
 - content source 133
 - creating kinetic content 44–45
 - creating stage 133
 - customer experience 75
 - definition of 3, 4, 8, 9, 20
 - delivery apps 135
 - formats 135
 - governance. *See* governance
 - governance role in 33–34
 - localization impact 123–124
 - management 133–134
 - metadata 128
 - of technical communication 3
 - operating models for 20
 - organizational
 - implementation 138
 - outlining and documenting content 127
 - personalized experiences 92–93
- robustness 43–44
- services 134
- structure and semantics 127
- supply chain 130–133
- technical goals 128
- technology 125
- content operations specialist 3
- content production 1, 4, 15, 20, 25.
 - See also* product content;
 - See also* creating kinetic content
- operating models for 41, 89
- content review and approval process 29
- content services 134, 142
- content source 133
- content strategies 4, 15
- content supply chain vs. content ecosystem 128
- content type 52, 83, 84, 85, 119, 122, 127
- content variants 28, 56
- ContentOps Manifesto 8
- continuous delivery pipelines 17, 45–46
- controlled natural language, writing technique 60–61
- conversation design 21
- conversion, customer experience 81
- copywriting 52, 53
- COVID-19 76, 82
- creating kinetic content 44–45
 - automate continuous delivery
 - pipelines 45–46
 - collaboration 45
 - content moves online 68–69
 - conventions and standards
 - usage 51–52
 - delivery 64–65
 - descriptive metadata 46
 - desktop publishing 67–68
 - editorial conventions 48–49, 52–54

highly structured
 ContentOps 71–72

information enablement 44

intelligent content 54–55

level of 66–67

metadata 46–47

multichannel delivery 70–71

people in creation
 process 149, 151

presentation opens doors 69–70

psychological signals 151

robustness 43–44

semantic information 46–47

service 64–65

structural metadata 46

technical capability 47–48

technical standards 49–51

work for people 153

writing techniques 59

creating stage 133

Current State of Component

- Content Management:
 An Integrative Literature Review, The (Andersen and Batova) 5

customer data platform (CDP) 93

customer decision journey 80

customer experience 75
 consideration 81
 conversion 81
 COVID-19 76
 customer journeys 77
 evaluation 82
 loyalty and advocacy 82
 marketing and digital 80
 omnichannel delivery 76
 perception (awareness) 81
 quantifying 88
 reengage 83
 return on investment (ROI) 78
 task-based customer journeys 80

customer journey map 77, 79

customer journeys 77
 brands 78
 content operations and 87–88
 designing 80
 map 77, 79
 omnichannel 85–86
 personalization and 86–87
 task-based 80, 85

customer life cycle 55, 83

D

Darwin Information Typing Architecture (DITA) 71, 124

DataOps 16

defining content standards 35

delivery apps 135

delivery pipeline 22

delivery, content 29–30, 64–65

descriptive metadata 46

Design for Cognitive Bias
 (Thomas) 149

design, progressive
 personalization 101–102

designing, personalization 105

DesignOps 16

DevOps 15

Dextrosinistral languages 49

digital asset management
 (DAM) 93

digital experience platform
 (DXP) 101

digital governance 33. *See also*
 governance

digital modeling 137

discoverable content 54

diversity, equity, and inclusion (DEI)
 63–64, 147, 148, 149, 150, 151

DocOps 16

double-entry bookkeeping 18–19

Dunbar, Natalie Marie 63

Dush, Lisa 6

dynamic content 65

E

Eagan, Colin 101, 101–102, 106
 Eagleson, Robert 59
 Ecma International 51
 e-commerce clothing company,
 content governance 36
 Edenic model 94
 editorial conventions 48–49, 52–54
 accessibility 62–63
 controlled language 60–61
 editorial structure to copy 57
 global readiness 61
 intelligent content 54–55
 operating model based on
 intelligent content 55–57
 plain language 59–60
 writing techniques 59
 editorial structure to copy 57
 minimalism 57
 modularity 58–59
 educational agenda, technical
 communicators 160–161
 error recognition and recovery 57
 evaluation, customer experience 82
 Evia, Carlos 155
Exponential Organizations
 (Ismail et al.,) 22

F

Facebook ads 152
 facial recognition algorithm 149
Filter Bubble, The (Pariser) 106
Forbes (Jenkins) 78
 formats 135
 freelancers 109, 113
 From Solo to Scaled (Dunbar) 63

G

GatherContent 9
Global Content Strategy
 (Swisher) 61

global readiness, writing
 technique 61
 globalization (G11n) 112
 Gollner, Joe 5, 64, 65
 good enough content 34–35
 governance 33
 and standards 34, 36
 checkpoints for 37
 defining content standards 35
 good enough content 34–35
 in content operations 33–34
 language 39
 planning 41
 post-production 38–39
 pre-production 36–37
 processes and tools 36–39
 production 37–38
 risk 39
 timeliness 39
 types of content 39
 urgency 39
 GPT (generative pre-trained
 transformer) 144–145. *See*
 grammar content 38, 92, 96,
 97, 102, 104, 106, 107,
 117, 122
 graph-based content 137

H

Hane, Carrie 4, 8, 70
 harnessing information
 enablement 45
Harvard Business Review
 (Gruner) 78
 Healthcare Information and
 Management Systems
 Society 45
 Horodyski, John 46

I

ICE (in-context exact) match 115

IEEE SA (the Institute of Electrical and Electronics Engineers Standards Association) 51
inclusion, writing techniques 63
inclusive writing 63
independent content 65
industrial field, technical communication 156–157
industry publishing tools 116
information enablement 44 harnessing information 45
leveraging and harnessing information 45
turning traditional functions 45
information management professionals 93
innovation, operating models 17, 18–19, 23
instructions for use (IFUs) 127
intelligent content 54–55 adaptable content 55
based on operating model 55–57 discoverable content 54
reconfigurable content 55 reusable content 54
Intelligent Content: A Primer (Rockley) 54
internationalization 110
internationalization (i18n) 112
Ismail, Salim 44
ISO (International Organization for Standardization) 50
issue/bug tracking tool. *See* query management

J

Jones, Colleen 8, 9, 66
JSON (JavaScript Object Notation) 72

K

key performance indicators (KPIs) 99, 100, 101

Kimball, Miles 6
kinesis 21, 44
kinetic content 20–22
knowledge management systems (KMS) 124

L

language localization 110
Latin script languages 49
Leading Content Design, (McConnell) 8
learning management systems (LMS) 124
leveraging and harnessing information 45
literate activity 66
localization 109 and globalization 112
client and supplier 112–113 content operations impact 123–124
globalization (G11n) 112
industry definitions 109, 111
industry publishing tools 116
internationalization (i18n) 112 language 110
machine translation (MT) 117–118
query management 118
technology elements 114 terminology management systems 116
translation management system (TMS) 115
translation memories 114–115
translation projects and process 118–121
working process 113
localization (l10n) 110
logographic languages 49
loyalty and advocacy 78, 82

M

machine translation (MT) 117–118

maintenance, content 28
 management process 133–134
Managing Enterprise Content: A Unified Content Strategy (Rockley) 14
 marketing automation platform (MAP) 93, 104
 marketing qualified leads (MQLs) 99
Mastering One Voice (Murdock and Randolph) 8
 maturity model 66, 67
 McConnell, Rachel 8, 66
 McKinsey & Company 150
 McKinsey (Diebner) 78
 mental model, governance 34
 metadata 5, 38, 53, 128
 administrative 46
 descriptive 46
 stages of production process 47
 structural 46, 55, 58
 Metadata Matters (Horodyski) 46
 Metts, Michael 63
 micro-journeys 80, 83
 migrate and upgrade content 142
 Miller, Carolyn 67
Minimal Manual, The (Carroll) 57
 minimalism 57
Minimalism beyond the Nurnberg Funnel (Carroll) 57
 modular content 58–59
 modularity 58–59
 molecular content 65
 multichannel delivery 29, 70–71, 76
 multi-language vendors (MLVs) 113
 Murdock, Toby 8

N

Neural MT (NMT). *See* machine translation (MT)

Nielsen Norman Group 16

Norman, Nielsen 16
Nurnberg Funnel, The (Carroll) 57

O

O’Keefe, Sarah 8, 9, 66, 67
 OASIS Open 51
 offered content 65, 129
 omnichannel customer experience 76, 81, 85–86
 operating models 14, 43
 activate strategy 16, 22
 based on intelligent content 55–57
 benefits 22–23
 case study 18–19
 collaboration 17, 22
 commonalities of 16–17
 content production 19
 continuous delivery pipeline 17
 DataOps 16
 delivery of information 17, 22
 DesignOps 16
 DevOps 15
 DocOps 16
 for kinetic content 20–22
 governance role in 33–34
 innovation 17, 18–19, 23
 operational objectives 17–18
 quality 18
 ResearchOps 16
 risk management 17, 23

Operating Models: Delivering on Strategy and Optimizing Processes (Murphy et al) 14

organizational implementation 138
 AI (artificial intelligence) 144–145
 battle card 139
 channels 139
 collaborators/constituents 140
 configure output/UX 144
 content 139–140
 GPT (generative pre-trained transformer) 144–145
 migrate and upgrade content 142

planning 141
 process of 141
 proven process 138
 stand-up and integrate
 technology 143
 technology selection 142
 training 143
 vendors 138
 outlining and documenting
 content 127
 ownership, content 36

P

page consistency 66
 pandemic-borne restrictions 151
 Pariser, Eli 106
 perception (awareness).
 See awareness
 performance, web page 38
 personae 80
 personalization 10, 91, 94. *See also*
 progressive personalization
 and customer journeys 86–87
 arguments for 92
 ContentOps advantage 92–93
 designing 91, 105
 digital interactions 95
 gap 93–99
 performance mindset 106–107
 precisionism 103–105
 prioritization 106
 progressive 96–99
 quintessential problem of 95
 regressive 101
 score performance 107
 simplification 105
 technology software 102
 zone-targeting approach
 102–103

Personalization Paradox, The
 (Swisher) 61

plain language, writing
 technique 59–60

planning 141
 governance 41
 point-of-sale (POS) interaction 81
 polylingual, centralized
 governance 40
 post-production governance 38–39
 precisionism, personalization
 103–105
 pre-production governance 36–37
 Pringle, Alan 66
 product content 21
 commonalities to 22–23
 product design 149
 product information management
 (PIM) 93, 104
 product life cycle 55, 134
 production governance 37–38
 profiled content 65
 progressive personalization
 96–99
 design 101–102
 solo project 99–100
 sonata project 100
 symphony project 100–101
 Project Inkblot 63
 project manager (PM) 113
 proprietary formats 136
 proven process 138
 publication management, technical
 communicators 158–159
 published content 129

Q

quantifying customer
 experience 88
 query management 118

R

radio-frequency identification
 (RFID) 50
 Randolph, Zoë 8
 reconfigurable content 55
 rendering, content 29

research operations 16
 ResearchOps 16
 return on investment (ROI) 78, 101
 reusable content 54
 review and approval process,
 content 29
 review, content 36
 rhetoric and stylistics, technical
 communicators 158
 risk management 17, 23, 31
 risk, content governance 39
 robustness, ContentOps 43–44
 content moves online 68–69
 desktop publishing 67–68
 highly structured
 ContentOps 71–72
 levels of 66–67
 multichannel delivery 70–71
 presentation opens doors 69–70
 Rockley, Ann 14, 54
 Rose, Robert 93

S

scalability 26–27
 scaling company activity 18
 score performance,
 personalization 107
 screenwriting format 49
 scriptwriting 49
 Searle, Loy 61
 segmentation models 80
 semantic information 46–47
 service, content 64–65
 Silver Surfer Adventurer
 persona 80
 Simplified Technical English 60
 single-language vendors
 (SLVs) 113
 sinistrodextral languages 49
smarter brands 82
 Social Justice Advocate Digital Native
 persona 80
 social media text and images 21

software as a
 service (SaaS) 100, 143
 solo project 99–100
 sonata project 100
 source content format 135
 digital modeling 137
 graph-based content 137
 proprietary formats 136
 special note 136
 text-based formats 136
 tree-based content 137
 special note format 136
 standard content processes 36–39,
 50–51
 stand-up and integrate
 technology 143
 Stanford and Cornell Universities
 case study 152
 Strategy& group of
 PricewaterhouseCoopers
 (PwC) 14
 structural metadata 46, 55, 58
 structure and semantics,
 content 127
 structured authoring 2, 66, 70, 158
 structured content 124
 subject matter experts (SMEs) 94
 suitability, content 36
 supply-chain-management 113,
 130–133
 support content 21
 Swisher, Val 61
 symphony project 100–101

T

task-based customer journeys 80,
 83, 84
 technical capability, content
 producers 47–48
 technical communicators/
 communication 6, 16, 47,
 69, 155
 academic/industry field 156–157

authentic connections 157
 educational agenda 160–161
 publication management
 158–159
 rhetoric and stylistics 158
 user experience and audience
 analysis 159–160
 technical content creators 25
 technical goals 128
 architecting confidence 129
 content supply chain vs. content
 ecosystem 128
 offered vs. published content 129
 technical standards 49–51
 technical writing 1, 3
Technical Writing 101 (O’Keefe) 66
 technical/human challenges 140
 technological affordance 67
 technological support 125
 technology selection 142
 technology, localization 114
 template-based authoring 66
 terminology management
 systems 116
 text-based formats 136
 Thomas, David Dylan 149
 timeliness, content governance 39
 training 143
 translation 109. *See also*
 localization
 projects and process 118–121
 video work/custom work 119
 translation (T9n) 110
 translation management system
 (TMS) 115
 translation memories (TMs)
 114–115
 translators 116, 118, 119, 121
 tree-based content 137
 turning traditional functions,
 information enablement 45

U

urgency, content governance 39
 user experience and audience
 analysis, technical
 communicators 159–160

V

van der Meij, Hans 57
 velocity 27
 authoring and editing 27–28
 delivery 29–30
 maintenance 28
 rendering 29
 review and approval process 29
 vendors 138
 video work/custom work,
 translation 119

W

W3C (World Wide Web
 Consortium) 51
Web Content Management
 (Barker) 13
 Welfle, Andy 63
 work for people 153
 writer glossaries 116
Writing Is Designing (*Metts and
 Welfle*) 63
 writing techniques 59
 accessibility 62–63
 controlled language 60–61
 diversity, equity, and inclusion
 (DEI) 63–64
 equity 63
 global readiness 61
 plain language 59–60

Z

zone-targeting model 101, 102, 106